Anterior Cingulate Cortex: Contributions to Social Cognition

Matthew Apps

August 2011

This thesis is submitted in partial fulfilment of the requirements for a degree of Doctor of Philosophy in Psychology at Royal Holloway, University of London.

Declaration of Authorship

I Matthew Apps hereby declare that this thesis and the work presented in it is entirely my own. Where I have consulted the work of others, this is always clearly stated.

Signed:

Date:

Abstract

It has been suggested that the Anterior Cingulate Cortex (ACC) plays an important role in decision-making. Activity in this area reflects processing related to two principles of Reinforcement Learning Theory (RLT): (i) signalling the predicted value of actions at the time they are instructed and (ii) signalling prediction errors at the time of the outcomes of actions. It has been suggested that neurons in the gyrus of the ACC (ACCg) process information about others' decisions and not one's own. An important aim of this thesis is to investigate whether the ACCg processes others' decisions in a manner that conforms to the principles of RLT. Four fMRI experiments investigate activity in the ACCg at the time of cues that signal either the predicted value of others' actions or that signal another's predictions are erroneous.

- Experiment 1: Activity in the ACCg occurred when the outcome of another's decision was unexpectedly positive.
- Experiment 2: Activity in the ACCg varied parametrically with the discrepancy between another's prediction of an outcome and the actual outcome known by the subject, in a manner that conformed to the computational principles of RLT.
- Experiment 3: Activity in the ACCg varied with the predicted value of a reward, discounted by the amount of effort required to obtain it.
- Experiment 4: Activity in the ACCg varied with the value of delayed rewards that were discounted in a manner that conformed to a social norm.

These results support the hypothesis that the ACCg processes the predicted value of others' actions and also signals when others' predictions about the value of their actions are erroneous, in a manner that conforms to the principles of RLT.

Acknowledgements

During the second year of my bachelor's degree, I was required to choose who I wished to supervise my final year project. Five and half years, one BSc. project, one MSc. project and now one PhD thesis later, I can say that choosing Dr. Narender Ramnani was one of the best decisions I have made (without any possibility of a 'prediction error' in the future!). He has been a wonderful mentor, guiding me along the path to the completion of this work and providing me with a level of training that I would not have received anywhere else. This is not to mention the considerable patience he has exercised with the most sarcastic and pushy of PhD students, who spent hour upon hour filling up his white board with incomprehensible algorithms. The success of this PhD and any future successes I may have in academia are a product of the amount time and quality of the time that he has given me. For all the opportunities he has extended to me, and for giving me the chance of pursuing research that I have thoroughly enjoyed, I offer my deepest gratitude. I would also like to thank my advisor Dr. Manos Tsakiris, who provided me with sound advice at several points during the PhD and by offering me future employment, provided me an important motivator to finish in less than three years.

I would also like to thank Dr. Mark Walton, Dr. James Kilner and Professor Johannes Zanker for all their time and effort they invested examining this thesis and for their thought provoking and useful comments they provided on the work at the Viva.

It is rare in life that you will meet someone that you can consider a good friend, who will also argue with you vociferously about how a computational model works. This highly unique position in my life has been filled by Elise Lesage. She helped me with programming ("rows, columns"), reminded me how numbers and formulas work, and made life in the office an enjoyable place to be. It is not an exaggeration to say that the work in this thesis would not have been of the same quality without her input. For everything she helped me with and for being such a dear friend, I am extremely grateful.

Having been at Royal Holloway for the best part of seven years, I have been lucky to make many friends and meet many wonderful people who have helped in numerous different ways. I would particularly like to thank Josh, who since my undergraduate days has offered me technical advice, wisdom on how to navigate the path of a PhD and extensive comic relief. For this and for providing comments on chapters, I offer my thanks. I would also like to thank Selina, Gary, Matt, Damian and Sam for providing useful comments on chapters, career advice, emotional support and being excellent company in a public house. I am particularly indebted

to Selina, who gave up a considerable amount of her time to keep me fed, watered, awake and sane during the last few late nights of writing this thesis. Also Sam, who made sure that the thesis was printed and bound in the right order and the right way up. I am also grateful to James, who has somehow managed to keep my overworked PCs running for three years and has always managed to do so with good humour.

Several people played crucial roles in the collection of data. I would like to thank Ari, for all his help with collecting MRI data. I also thank my confederates, Robin, Erman, Eden, Javier, Matt, Arleta and Vivien. Their ability to lie convincingly was a crucial component of the work in this thesis. Eden must also be thanked for helping with the creation of cingulate masks used in chapter four.

When I began this PhD three years ago, Robina did not realise that this meant that she would also be doing a part-time, unfunded PhD herself. She has shown immeasurable patience over this time; attending a conference, reading through several chapters of the thesis and listening to my ramblings about the importance of a few millimetres of grey matter to social behaviour. I have been extremely lucky to have her there with me throughout the PhD, supporting me with love, kindness and affection.

Finally, I would like to thank my family. I am extremely fortunate to have such loving, patient and good humoured family members. I do not know how I would have completed this work without you.

Table of Contents

Table of figures

Table of tables

Chapter 1: Introduction

1.1 Introduction

For many years it has been known that damage to the frontal lobe can affect social behaviour. The famous case of Phineas Gage highlighted how damage to the frontal lobe can change many aspects of an individual's personality, including how they behave during social interactions. More recent evidence has suggested that damage to the medial wall, particularly to the Orbitofrontal Cortex (OFC) and the Anterior Cingulate Cortex (ACC), can result in patients being afflicted with severe deficits in social behaviour (Tranel et al., 2002). Indeed, some have suggested that patients with such damage suffer from "acquired sociopathy" (Saver and Damasio, 1991), whereby the brain damaged patients behave similarly to psychopathic patients. However, the most prominent theoretical accounts of social cognition have not considered the ACC or the OFC as regions which are important for social cognitive abilities (Gallese and Goldman, 1998; Frith and Frith, 2006; Keysers and Gazzola, 2007; Schubotz, 2007). In this thesis, it will be argued that a portion of the ACC plays an important role in social cognition.

Experimental lesion studies, which have the advantage of greater anatomical specificity in the location of the lesion compared to patients with acquired damage, have suggested that disruption of the ACC and not the OFC are the cause of the perturbed social behaviour (Hadland et al., 2003). More specifically, lesions to the gyral surface of the Anterior Cingulate Cortex (ACCg) and not the sulcus (ACCs) in the macaque monkey, have been shown to disrupt normal social behaviour and also the processing of social stimuli (Rudebeck et al., 2006a). Interestingly, these two regions have distinct anatomical properties. The cytoarchitecture of the gyral surface and the sulcus are distinct along the entire extent of the ACC (Vogt et al., 1995). Similarly, there are distinctions in the connectional properties of these two regions; with the ACCg connected to areas of the brain that are implicated in processing social information that the ACCs is not connected to.

How does the ACCg contribute to social cognition? Despite the evidence implicating the ACCg in processing social information, there has been a notable absence of any theoretical accounts of the role of this region in social cognition. However, one seminal fMRI study in humans highlighted how the ACCg may process similar information to the ACCs, particularly during decision-making. However, the ACCs processes the information that guides first-person decisions, whereas the ACCg processes this information about others' decisions (Behrens et al., 2008). To understand how the ACCg contributes to social cognition it is therefore important to understand the contribution of the ACCs to one's own cognitive processes. Whilst there is an absence of a literature investigating what socially relevant information is processed in the ACCg, there are a considerable number of studies that have investigated the functional properties of the ACCs.

What information is processed in the ACCs? Several theoretical accounts of the functional properties of the ACCs have been proposed (Bush et al., 2000; Holroyd and Coles, 2002; Botvinick, 2007). However, these theories cannot account for all of the data that speaks to the functional properties of the ACCs. In this chapter, an alternative framework for information processing in the ACCs will be outlined. This framework is based around two important notions; firstly, that the ACCs processes how rewarding an action is, discounted by other costs associated with performing the action in order to receive the reward, and secondly, that information processing in the ACCs conforms to the principles of a well established model of learning, namely Reinforcement Learning Theory (RLT). The two most important components of RLT are; firstly that predictions are made about the outcomes of actions that guide decisions between different actions, and secondly, that these predictions are updated when new information reveals the prediction was erroneous. In this chapter, lesion studies, neurophysiological recordings and functional imaging evidence will be cited that supports the claims that the ACCs processes predictions about the discounted value of actions and also signals when predictions are erroneous. In this thesis, it will be argued that this framework can also be applied to the ACCg, to explain how this area processes social information. Specifically, this thesis makes the claim that the ACCg processes the variables that guide other's decision-making and processes them in a manner that conforms to the computational principles of RLT.

In this specific chapter the anatomical and functional properties of the ACCg will be outlined, highlighting the area as a candidate for processing social information. I will then argue that the ACCg processes similar information to the ACCs. Therefore, in order to generate hypotheses about the contribution of the ACCg to social cognition, a discussion of the theoretical accounts of ACCs function will be provided, before a framework for how the ACCs contributes to first-

person decision-making is outlined. Each of the four studies in this thesis will examine whether the ACCg codes the same information as the ACCs, but processes this information about the actions and decisions of others. Specifically, two questions are asked:

(i) Does the ACCg process information in a manner that conforms to RLT?

(ii) Does the ACCg process information about the variables that guide others' decisions?

1.2 The Social Brain

Many species live in complex social environments, in which survival depends on the ability to interact successfully with other conspecifics. In their lifetime each individual will engage in a large number of social interactions that can take a variety of different forms. In addition, for each social exchange to be successful, a diverse array of complex cognitive processes must be performed. The significance of these social cognitive abilities is underscored by the link between forebrain size, the number of conspecifics with which an individual interacts and the complexity of the interactions in which an individual engages (Dunbar and Shultz, 2007; Shultz and Dunbar, 2007; Shultz and Dunbar, 2010). Increases in the size of a social group and the complexity of interactions with individual members of that group, are correlated with interspecies increases in the brain size of birds and several mammalian species including humans (Silk et al., 2003; Holekamp et al., 2007; Silk, 2007). Social cognitive abilities are therefore considered as one of many important selection pressures that are linked to an individual's survival and evolutionary fitness (Dunbar and Shultz, 2007; Silk, 2007). Identifying and understanding the neural mechanisms that underpin social cognitive processes has therefore become a hugely important area of research.

It is unequivocal that humans engage in the most complex social behaviour of any species. This complexity has meant that a comprehensive understanding of the cognitive processes that underpin social behaviour has not yet been achieved. As a result, localizing the neural antecedents of social cognitive abilities has been equally challenging. Historically, there were two accounts of the cognitive processes that underpinned social behaviour. One account suggests that one's own cognitive processes are mirrored when trying to understand the same cognitive processes of another (Gallese and Goldman, 1998). As such, individuals understand the intentions and actions of others by simulating them, as if having the same intentions or performing the same actions themselves (i.e. "putting oneself in their shoes"). Evidence to support the claims of simulation theory has been provided by the discovery of mirror-neurons in the macaque premotor and parietal cortices (Gallese and Goldman, 1998; Rizzolatti and Craighero, 2004). The defining property of these neurons is that they increase their spike rate when a monkey makes a goal-directed action and also when the same actions are observed being performed by another. An alternative account of social cognition suggests that we understand others by making theories about the desires, attitudes and beliefs that drive their behaviour (Baron-Cohen et al., 1985). The ability to understand and represent the mental

states of others is often referred to as “theory of mind” or “mentalizing”. This account makes the specific prediction that the processing of others’ mental states will occur in networks that are specialised for the functions necessary to deduce the intentions, desires and beliefs from others’ behaviours. Support for this claim has come from neuroimaging studies in humans that consistently find activity in three areas when subjects are processing the mental states of others. This core-circuit of areas which is engaged when mentalizing consists of the paracingulate cortex (a region lying at the border of areas 8, 9 and 32’ on the medial wall), a region in the depths of the posterior Superior Temporal Sulcus (pSTS) lying adjacent to the Tempero-Parietal Junction (TPJ) and a region at the tip of the temporal poles (Frith and Frith, 2003; Frith and Frith, 2006).

In this section I review some of the evidence that supports each theory. I will then highlight how neither can account for how damage to the ACC, an area which is not considered as part of the core-circuit or the mirror-neuron system, can significantly disrupt social behaviour.

1.2.1 Simulation Theory

Simulation theory states that we understand the actions and goals of others by simulating them in circuits that process our own goals and actions (Gallese and Goldman, 1998). The first evidence to support this notion came from the discovery mirror-neurons that increase their spike rate when a monkey makes a goal-directed action and also when the same actions are observed being performed by another. Neurons that have such properties were first found in area F5 in the macaque monkey premotor cortex (Gallese et al., 1996; Rizzolatti et al., 1996). Following this, neurons with similar properties were identified in area PF in the macaque parietal lobe (Fogassi et al., 2005). Neuroimaging studies have also found that homologous areas of the human brain have similar mirror-like properties, activating both when observing an action, or performing the same action oneself (Buccino et al., 2001; Buccino et al., 2004; Iacoboni et al., 2005). The notion that the premotor and parietal cortices have mirror-like properties is supported by an fMRI study by Buccino et al., (2001). Subjects observed one of three actions, the biting of an apple (a mouth movement), reaching and grasping of an object (a hand movement) and the kicking of a ball (a foot movement). Activity was compared between these conditions and conditions where the same body part was viewed when static. They reported activity in both parietal (area 7) and premotor areas (BA 44) that are homologous to area PF and area F5 in the macaque brain respectively, when subjects observed

the actions of another. Interestingly, they also found that the activity at the time of these actions was organised topographically in the same manner that the premotor cortex responds in a topographic manner when performing an action oneself, i.e. the area that responded to the foot actions was located medially, whereas the area that responded to the mouth actions was lateral. Other studies have shown that area 44 is also engaged when others' actions can be predicted. Ramnani and Miall (2004) performed an fMRI study in which subjects monitored cues that instructed a third-person or a computer to perform an action, and also performed similarly cued actions themselves. When an abstract cue informed the subject of the specific action that would be taken by the third-person, the premotor cortex was activated. When similar cues were uninformative about the specific action that would be taken by a third-person, the premotor cortex was not activated. They also reported that a portion of the premotor cortex was activated when cues instructed the specific action that they would be required to perform. These results highlighted the mirror-like properties of human premotor cortex. Numerous fMRI studies have since supported the assertion that both the parietal lobe and the premotor cortex have mirror-like properties, although there is still considerable debate as to whether they can be considered a "mirror-neuron" system in humans in the same way as they are in the monkey (Rizzolatti and Craighero, 2004; Iacoboni and Dapretto, 2006; Kilner et al., 2009; Kilner, 2011). However, these studies highlight how recognising others' actions, requires the ability to simulate those actions in functional anatomy that guides one's own actions.

1.2.2 Theory of Mind

Another prominent account of social cognition suggests that we understand the behaviour of others by making theories about their mental states. Since the development of neuroimaging methods, a large number of studies have investigated the neural processes that underpin such theory of mind abilities. A striking finding is that across a variety of different tasks, in a large number of different studies, three areas are consistently reported as activated, namely the paracingulate cortex, the posterior Superior Temporal Sulcus (pSTS) and the temporal poles. This includes studies that have required subjects to read stories (Fletcher et al., 1995; Gallagher et al., 2000), read cartoons (Gallagher et al., 2000; Sommer et al., 2007), monitor animations (Castelli et al., 2000) or monitor others' actions and decisions (Ramnani and Miall, 2004; Rilling et al., 2004). The consistent activation in these three areas across studies and the

fact they are anatomically connected to each other (Seltzer and Pandya, 1989; Barbas et al., 1999) has led to them being referred to as the core-circuit for mentalizing.

The region of the paracingulate cortex that is typically activated by mentalizing tasks lies superior to the cingulate sulcus, putatively at the borders between areas 32', 8 and 9 (Amodio and Frith, 2006). Whilst this area has been found to be activated in many tasks where mental states have to be attributed to others, there is some evidence to suggest that this area has a more specific role. It has been claimed that the paracingulate cortex is engaged specifically when the intentions that another wishes to communicate are inferred (Amodio and Frith, 2006). One study that supports this claim was conducted by Walter et al. (2004). They reported activity in the paracingulate cortex when subjects saw cartoons in which the intentions of another to communicate during a social interaction could be inferred. Activity was not found when a mental state could be attributed to another, but the actions of another were not communicative (Walter et al., 2004). Other studies have since supported the claim that the paracingulate cortex is important for the processing of communicative intentions (Walter et al., 2009).

There is still considerable debate as to the contribution of the pSTS to mentalizing abilities, however it appears to have a specific role in detecting that motion in a scene is caused by a biological agent. For example, neurons in the pSTS increase their firing rate when the motion of point lights takes the form of a moving biological agent, which is not the case when the movement of the point lights is random (Puce and Perrett, 2003). Neuroimaging studies have supported this claim, showing that the movement of limbs and also faces moving during speech, result in an increased BOLD response in the pSTS (Puce et al., 1998; Puce and Perrett, 2003; Pelphrey et al., 2004; Thompson et al., 2007). Other studies have shown that activity occurs in this area even when abstract objects interact in a manner that suggests the movement is of two biological agents interacting (Castelli et al., 2000). Some have therefore suggested that the role of pSTS may be to code for instances when the motion detected in a scene is that of another organism (Frith and Frith, 2003) .

The polar region of superior temporal gyrus, lying adjacent to the amygdala, is often activated in tasks in which subjects attribute mental states to others (Frith and Frith, 2006). This would implicate the pSTS in processing information about the mental states of others. However, there is some evidence to suggest that processing in this area is not actually related to mental states. There are a number of studies that have shown that this area is engaged when processing semantic information. When retrieving the meaning of words during sentences, the

temporal pole shows an increased BOLD response (Mazoyer et al., 1993; Fletcher et al., 1995). This area is also activated when performing lexical decision-making tasks and when making semantic judgements (Noppeney and Price, 2002). This would suggest that the temporal poles may be engaged by the retrieval of semantic information and not specifically when processing the mental states of others. Some support from this comes from the fact that the majority of tasks that examine mentalizing processes which report activity in the temporal pole are those which use vignettes or cartoon based stimuli (Fletcher et al., 1995; Gallagher et al., 2000; Vogeley et al., 2001; Frith and Frith, 2003). Studies that use tasks where the stimuli do not contain any semantic information, do not show activity in the temporal poles when subjects are attributing a mental state to others (Castelli et al., 2000; Ramnani and Miall, 2004; Rilling et al., 2004). Thus, there is no clear picture as to whether a portion of the temporal poles is engaged specifically when processing others' mental states. However, in general, there is considerable evidence to suggest that the pSTS and the paracingulate cortex form a network that processes others' mental states.

1.2.3 Theory of Mind, Simulation and the Anterior Cingulate Cortex

It is clear from the evidence provided that parts of the motor system are activated both when observing and performing the same action. In addition, when processing others' intentions, desires and beliefs there is activity in a core-circuit of areas that may be specialised for such processes. However, given the complexity of social behaviour, an obvious question arises: Can these theories explain all aspects of social behaviour and social cognition? In this thesis, it will be argued that these theories, in their current form, cannot account for how we understand the actions and decisions of others. In this next section, I provide evidence which suggests that the integrity of areas that are not part of the mentalizing system or the mirror-neuron system may also be important for social cognition.

Neurological patients offered the first insight into how areas on the medial wall in the frontal lobe (inferior and rostral to the paracingulate cortex) might contribute to social cognition. Patients with lesions to this region often show changes to their social behaviour in comparison to their behaviour before a lesion, without any changes in performance on other cognitive tasks (Barrash et al., 2000; Tranel et al., 2002; Bar-On et al., 2003). A series of patients who have suffered significant damage particularly to the medial wall including the OFC, the ACC and ventral portions of the superior frontal gyrus, have shown that changes in social behaviour can

be quite severe (Saver and Damasio, 1991; Barrash et al., 2000). Indeed, some have suggested that such damage can result in "acquired sociopathy" (Saver and Damasio, 1991), whereby patients respond to emotional and social stimuli in the same manner as psychopathic patients. Such symptoms are evident in patient E.V.R, who had extensive damage across a large portion of the ventral areas of the medial wall (Saver and Damasio, 1991; Anderson et al., 1999). When patient E.V.R observes distressing social situations, the change in skin conductance that are observed in the normal population do not occur. This absence of a change in skin conductance is also present in patients with psychopathy (Anderson et al., 1999).

Whilst it is difficult to impute function to a specific brain area following acquired damage which is not constrained to a specific anatomical location, there is some evidence that it is damage to the ACC that results in changes in social behaviour. Tranel and colleagues (Tranel et al., 2002) examined what they termed 'social conduct' in seven patients who had damage to the medial wall in ventral portions of the frontal lobe. They examined the changes that occurred in patients' social behaviours following the lesion, as measured by a clinical assessment, their record of employment, their social status and their interpersonal skills (measured by reports from close relatives). Four patients were rated as showing significant changes and deficits on each of those four measures. In their paper, Tranel and colleagues (2002) reported that each of these patients had lesions that extended predominantly into the right hemisphere and not the left hemisphere, arguing that social cognition may be lateralized. However, it is notable that these patients all had lesions that extended over large portions of the ACC, whereas the patients with lesions that were constrained to the OFC or the most subgenual portions of the ACC showed very little change in their social behaviour. This suggests that it could be lesions to the ACC in these patients that disrupts their social cognitive abilities, not lesions to the OFC or damage in one hemisphere. Thus, the tentative conclusion that can be drawn from this study that it is lesions to the ACC and not the OFC that cause deficits in social behaviour.

Patients who have undergone a cingulotomy (removal of grey matter in the cingulate cortex) offer the possibility of examining how damage to the ACC changes behaviour. A cingulotomy typically involves lesions to the dorsal portions of the ACC, to relieve symptoms of chronic pain or Obsessive-Compulsive Disorder (OCD) (Ballantine et al., 1967; Ballantine et al., 1987; Jenike et al., 1991; Wilkinson et al., 1999). A number of studies have reported behavioural and cognitive changes that are consequent from a bilateral cingulotomy, beyond the improvement in the symptoms of the disorder that was being treated by the lesion. Typically these studies have shown that a cingulotomy causes deficits in generating actions and sustaining attention

during cognitive tasks (Janer and Pardo, 1991; Cohen et al., 1999). However, there is little evidence that a cingulotomy results in deficits in social behaviour (Janer and Pardo, 1991; Jenike et al., 1991; Cohen et al., 1999; Cohen et al., 2001; Ochsner et al., 2001; Davis et al., 2005). Indeed, the only evidence of changes in social behaviour following a cingulotomy, are that many patients report an improvement in relations with others in the months after the surgery (Wilkinson et al., 1999; Cohen et al., 2001). However, this is most likely due to the improvement in the symptoms of the disorder for which the surgery was conducted.

One important point to note is that the studies that have investigated behavioural deficits following a cingulotomy, have all investigated deficits in patients with lesions to the sulcus of the ACC and not the gyrus. This is not surprising, as a cingulotomy typically involves lesions to the sulcus and not the gyrus. In the next two sections of this chapter I will report evidence that suggests it is the gyrus of the ACC that is important for social behaviour. This may explain why there is an absence of any reported deficits in social behaviour in patients who have undergone a cingulotomy.

This section has outlined evidence which suggests that the ACC an area of the brain that is not considered part of a Theory of Mind network or the mirror-neuron system, is important for human social cognition. In later sections, I will argue that it is a specific portion of this region, on its gyral surface, that is engaged when processing social information. However, before there is any further discussion of the functional properties of the area, the anatomical properties of the ACC will be discussed, so that future discussions about the localization of functions within different portions of the ACC can be made coherently

1.3 Anatomy of the Anterior Cingulate Cortex

The aim of cognitive neuroscience over the last 60 years has been to localize function in the brain. Few cortical areas have been ascribed more functions than the Anterior Cingulate Cortex (ACC). Functional imaging, neurophysiological and electrophysiological recordings, electrical stimulation and lesion studies suggest that the ACC has diverse functional properties. Some of

this research has implicated the ACC in autonomic processes, including the regulation of blood pressure, heart rate and breathing (Ward, 1948; Kaada et al., 1949; Pool and Ransohoff, 1949; Showers, 1959; Terreberry and Neafsey, 1984, 1987; Liotti et al., 2001). Other studies report this area as being important for the control of limb movements (Luppino et al., 1991). One might conclude from this that the ACC is involved in the control of low-level, predominantly visceral processes. However, seemingly contradictory evidence highlights the ACC as important for higher-order cognitive and emotional processes, including volition, self-awareness and even conscious experience (Devinsky et al., 1995; Posner and Rothbart, 1998; Bush et al., 2000; Allman et al., 2005; Posner et al., 2007). One could try and reconcile these two viewpoints by identifying the common property that underpins each of these processes. However, an alternative approach might be to examine the validity of the assumptions that such research makes.

Much of the research investigating the functional properties of the ACC, particularly those that use neuroimaging methods, are underpinned by the assumption that the ACC is an anatomically homogenous zone. However, examination of the anatomical properties of the region often referred to as ACC, invalidates such an assumption. Neither the cytoarchitecture nor the afferent or efferent connections are homogenous across the spatial extent of the ACC. There is also evidence to suggest that many of the different functions that are processed in the ACC can be localized to distinct sub-regions which have different anatomical properties (Beckmann et al., 2009). In this chapter, it will be argued that the processing of social information can be localized to a portion of the gyral surface of the ACC. In this specific section, the anatomical properties of the cingulate cortex will be outlined, focusing on the differences in cytoarchitecture and connectional properties of the Anterior Cingulate Sulcus (ACCs) and the Anterior Cingulate Gyrus (ACCg).

1.3.1 Cingulate Cytoarchitecture

"The functions of a cortical area are determined by its extrinsic connections and intrinsic properties"

(Passingham, Stephan & Kotter, 2002)

An important notion in cognitive neuroscience is that understanding the structure of the brain is imperative for understanding its function. Since the work of Brodmann (1909) it has been well established that the brain can be anatomically divided into separate regions based on its intrinsic cytoarchitecture (Brodmann, 1909). Such divisions have proven extremely useful for functional localization and also for examining homologies between species (Passingham et al., 2002). In this section I will outline the cytoarchitectonic properties of the cingulate cortex. The evidence provided will support the case that the sulcus and the gyrus have different functional properties.

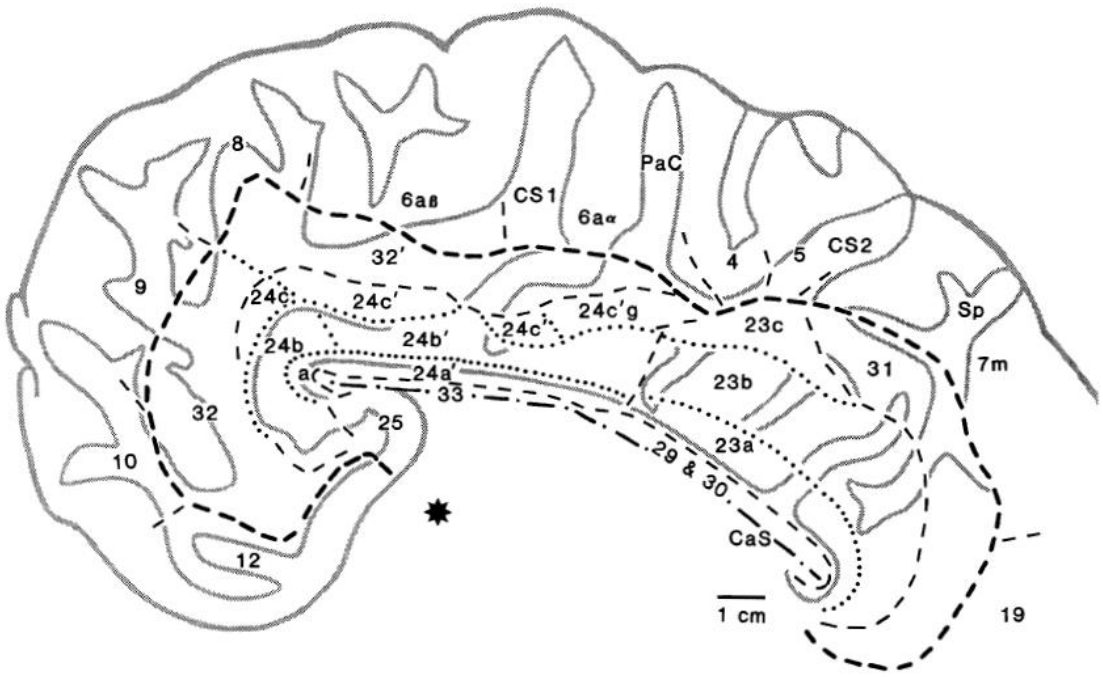

*Fig.1.1. Illustration of human cytoarchitecture on the medial wall taken from (Vogt et al., 1995). The frontal lobe is shown to the left. The * shows the location of the anterior commisure. The cingulate cortex encompasses all the regions within the bold dotted line.*

The human cingulate cortex extends along the medial wall of both hemispheres (see fig.1.1) around the entire extent of the corpus callosum. The anterior and rostral portion extends around the genus ('subgenually') and the posterior portion extends around the splenium. The cingulate gyrus is segregated from the parietal lobe by the marginal ramus of the cingulate

sulcus. This sulcus extends rostrally segregating the cingulate gyrus from areas 6, 8, 9 and 10 in the frontal lobe (Vogt et al., 1995; Paus et al., 1996; Paus, 2001; Palomero-Gallagher et al., 2008). Whilst cingulate sulci are present in all subjects and in all hemispheres, there is considerable variability in terms of the presence and location of an additional paracingulate sulcus or superior cingulate sulcus (Vogt et al., 1995). Paus et al., (1996) examined high resolution Magnetic Resonance Imaging (MRI) scans of 494 human hemispheres (297 brains). A paracingulate sulcus was present in 88% of cases overall. However, it was only prominent in 45% of brains with a slightly higher incidence in the left hemisphere.

Possibly the most detailed account of human and monkey cytoarchitecture has been provided by the work of Vogt and his colleagues (Vogt et al., 1987; Vogt et al., 1995; Palomero-Gallagher et al., 2008). Although there are alternative characterisations of the cingulate cytoarchitecture (Ongur et al., 2003), the work of Vogt will be used throughout this thesis. Broadly speaking the human cingulate cortex can be divided into four areas, retrosplenial, posterior (PCC), mid (MCC) and anterior cingulate cortices (ACC), although there are no gross anatomical landmarks to delineate them. The border between the PCC and the MCC lies on average, 22mm posterior to a vertical plane in line with the Anterior Commisure (VCA). The border between MCC and ACC lies, on average, 30mm anterior to the VCA, with the ACC extending around the genus of the corpus callosum (Vogt et al., 1995). However, these borders are only rough estimates, due to the considerable variability across subjects in their location and also as they do not lie perpendicular to the VCA line (Vogt et al., 1995; Palomero-Gallagher et al., 2008). It is important to note that both the anatomically defined ACC and the MCC fall in the region that is typically labelled as ACC by functional imaging research. I will therefore restrict further discussion to the cellular properties of these areas.

Both the MCC and the ACC can be further sub-divided into separate zones (Vogt et al., 1995). In the most rostral portion of the ACC, in subgenual cortex is area 25. Area 25 is defined by broad layers II and III, which are poorly differentiated. Layer V and VI are also poorly differentiated with large and densely packed neurons. Area 32 borders area 25, extending in the anterior plane to it. Area 32 sweeps caudally on the upper bank the of cingulate sulcus and is a dysgranular zone between agranular area 24c/24c' and the granular frontal areas of 9 and 10. In cases where there is an additional paracingulate sulcus, this area is typically lies within the depths of this additional sulcus (Vogt et al., 1995; Bozkurt et al., 2005). It has a broad layer II, with a well differentiated layer III containing large pyramidal cells. It also has a well differentiated layer IV containing some layer Va pyramids, which is typical of dysgranular cortex.

Below area 32 in the lower bank of the cingulate sulcus is area 24c, although in cases where there is an additional paracingulate suclus area 24c falls on both upper and lower banks of the cingulate sulcus. 24c also extends caudally from a border with area 25. Area 24c lacks a granular layer IV, but has a cell-dense layer Va and broad layer II, with no layer Vb, making it distinct from area 32. Layer III contains medium sized uniformly distributed pyramids. Area 24b, lies adjacent to area 24c extending caudally from a border with area 25. However, area 24b is found predominantly in the gyrus, even in cases where there is an additional paracingulate sulcus. This area also lacks a layer IV, but has a highly densely packed layer Va, containing both small and large pyramids. Area 24a lies adjacent to area 24b, sweeping caudally from its border with area 25. Area 24a is found solely on the gyral surface and does not encroach upon the sulcus at any point. Like areas 24c and 24b, 24a has no granular layer IV, with well differentiated layers II and III. However, it has a much thinner layer Va than area 24b. Finally, area 33, which lies in the depths of the cingulate cortex, lies in the callosal sulcus, adjacent to the corpus callosum. This area extends across both ACC and MCC. Area 33 is the most poorly differentiated area in the cingulate cortex. Layers II and III are broad and undifferentiated, layer V has sparse pyramidal neurons and there is no layer VI.

Lying caudal and posterior to ACC areas 32, 24a, 24b and 24c are areas 32', 24a', 24b' and 24c' in the MCC (Vogt et al., 1987; Vogt et al., 1995). The use of the same numerals, despite a transition from ACC to MCC, is unfortunately somewhat misleading. There are clear distinctions in the properties of these areas. Area 32' lies caudal to area 32 on the dorsal bank of the cingulate sulcus, but when an additional paracingulate sulcus is present it falls within this additional paracingulate sulcus and not in the cingulate sulcus. Area 32' has an attenuated layer IV with respect to area 32, its layer Va is less dense and layer III contains larger pyramidal neurons. Layer 24c' lies caudal to area 24c on the dorsal bank of the sulcus in the MCC. When there is an additional paracingulate sulcus, area 24c' is found on both the dorsal and ventral banks of the cingulate sulcus. Layer 24c' is distinct from area 24c as it contains medium sized pyramidals in layer Vb, a layer which is absent in area 24c. In addition, layer III is much more sparse in area 24c' than area 24c. Area 24c' is also distinct from area 32', as it does not contain the large pyramidal neurons that's form a broad layer III and IIIc in area 32'. Area 24b' lies below area 24c' and caudal to area 24b. Area 24b' has an attenuated Layer III compared to area 24b, has a thinner area layer Va and has clusters of large pyramidal neurons. However, layer Va is more dense in area 24b' than in area 24c'. Finally, layer 24a' lies below area 24b' on the gyral surface and caudal to area 24a. Area 24a' has better differentiated layers II and III

than in either areas 24a or 24b'. Layer Va is also thinner in area 24a' than in 24a (Vogt et al., 1995; Palomero-Gallagher et al., 2008).

In summary within both the ACC and the MCC there is a clear distinction between the cytoarchitecture in the cingulate gyrus and in the cingulate sulcus (Vogt et al., 1995). This is particularly apparent in the MCC, where broadly speaking, areas 24c' and 32' lie on the ventral and dorsal banks of the sulcus respectively and areas 24a' and 24b' lie on the gyral surface. This anatomical evidence therefore supports the claim that the sulcus and the gyrus of the cingulate cortex are engaged during different processes. In later sections I will argue that it is the gyral portion of the MCC that processes social information.

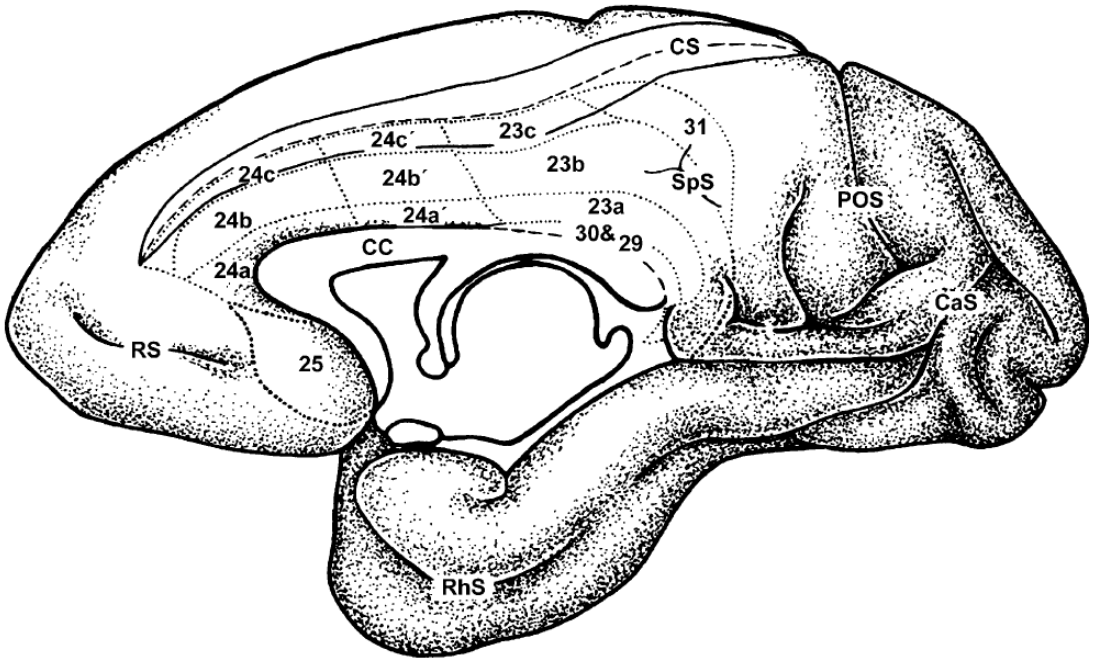

Fig.1.2. Illustration of the macaque monkey cytoarchitecture of the medial wall. Taken from (Vogt et al., 1987).

During this thesis I will often use evidence from lesion studies and neurophysiological recordings in non-human primates to make claims about how function is localized in the cingulate cortex. In particular, I will often refer to work that has been conducted on macaque monkeys. It is therefore important to discuss whether there are homologies between the cingulate cortex of the macaque monkey and that of the human. Figure 1.2 gives a clear illustration of the cytoarchitecture of the macaque monkey on the medial wall (Vogt et al., 1987). As can be seen, the macaque cingulate cortex can also be segregated into ACC, MCC

and PCC regions as in the human. Vogt and colleagues (Vogt and Pandya, 1987; Vogt et al., 1995) also claim that the each region can be subdivided into the same divisions as those found in the human, using the same criteria. The macaque cingulate cortex also contains a similar rostro-caudal differentiation in cytoarchitecture as that of the human. As in the human, area 25 is located subgenually and areas 24a-c lie caudal to area 25, but rostral to areas 24a'-c'. However, it is notable that the cingulate sulcus and gyrus do not extend around the genus of the corpus callosum to the same extent as the human cingulate sulcus and gyrus.

The cytoarchitecture of the macaque also shows a similar differentiation between the sulcus and the gyrus. Areas 24c and 24c' lie predominantly on the ventral bank of the sulcus, with areas 24a, 24a' 24b and 24b' inferior to the sulcus on the gyral surface. Thus, there is a considerable homology between the cingulate cortex in the macaque and the human brain (Vogt et al., 1987; Vogt et al., 1995). Despite the homologies, the cellular properties of the cingulate cortex have evolved. Historically, many regarded the cingulate cortex as a phylogenetically primitive region, due to its composition of only five rather than six layers. However, more recently the opposite claim has been made. This reverse has come about through the discovery of a class of large spindle neuron (sometimes called " Von economo neurons") which are present in areas 24b, 24b', 24a' and some have argued in area 24c' (Vogt et al., 1995; Nimchinsky et al., 1999). These neurons are found exclusively in the brains of humans, great apes, elephants and some cetacean species, but are greatest in number and density in humans (Nimchinsky et al., 1999; Allman et al., 2002; Watson et al., 2006; Butti et al., 2009; Allman et al., 2010; Allman et al., 2011). However, as yet there is little evidence of what specific functional properties the presence of these neurons imputes to a region

This evidence suggests that distinct functions may be processed in the ACC and the MCC, and moreover, different information may be processed within their respective subdivisions. However, such information is not particularly useful for understanding the functional properties of each subdivision alone. Despite some cellular properties linking closely to the function of a region (e.g. giant pyramidal cells in deep cortical layers typically implicate a role in motor control) the cellular properties of an area alone are not enough to impute a functional property. Thus, to understand the function of an area one must look beyond its intrinsic properties, to its extrinsic connections to other areas (Passingham et al., 2002).

1.3.2 Connectional anatomy of the Anterior Cingulate cortex.

The location and strength of the anatomical connections of an area can reveal important clues about the information it processes (Passingham et al., 2002). In this section I will review studies that have investigated the anatomical connections in the nonhuman primate and human cingulate cortex. Broadly speaking, both species show similarities in how the cingulate cortex can be divided into distinct areas that have different connectivity profiles. In both species the connections of the cingulate gyrus are, at least partially, distinct from the connections of the sulcus. Moreover, the gyrus in the MCC has connections to areas of the brain that are implicated in processing social information, that are not evident in other portions of the cingulate cortex.

1.3.2.1 Animal studies

Tracer studies have been used to investigate anatomical connections in the cingulate cortex. Within the limits of practicality, (i.e. constraining the injection of tracer material to a cytoarchitectonic zone is not always possible) it is possible to use tracers and cytoarchitecture to determine the correspondence between connectivity and cellular architecture. In the cingulate cortex, tracer studies have revealed that the ACC and the MCC share connections to areas of association cortex and areas that process rewards. However, the sulcus of the MCC is the only region of either the MCC or the ACC that has strong connections to the motor system. In addition, the gyrus of the MCC is the only region of either the MCC or the ACC that has connections to each of the nodes in the core-circuit that is engaged when mentalizing.

All of the areas in the ACC (25, 24a and 24b) have reciprocal connections with areas 46 and portions of area 9, both on the medial surface and also more lateral portions (Pandya et al., 1981; Vogt and Pandya, 1987; Petrides and Pandya, 2006, 2007). These areas are considered as association cortex and process highly abstract motor information (Petrides and Pandya, 1999). Subgenual portions of the ACC (area 25) have connections to TS2 and TS3 in the temporal lobe which are considered to be portions of the auditory association cortex (Friederici, 2002), although only weak connections are found from the more caudal portions of the ACC (areas 24a and 24b). Perhaps the strongest connections from the subgenual ACC are to the adjacent orbitofrontal cortex (OFC) (Morecraft et al., 1992; Carmichael and Price, 1995; Cavada et al., 2000). In the more subgenual regions these connections are primarily to medial areas 11 and 14, whereas as the strongest connection in caudal portions of the ACC are found to the more

lateral portions of area 11, as well areas 13 and 12 in the lateral OFC. Connections between the caudal ACC area and the medial OFC are present, but weaker than in the subgenual area (Morecraft et al., 1992; Carmichael and Price, 1995; Cavada et al., 2000). The OFC is well known for its role in processing rewarding stimuli (Schultz et al., 2000), with neurons in this area sensitive to the magnitude of reinforcers (Padoa-Schioppa and Assad, 2006, 2008). All of the areas in the ACC also project to another area which is considered as important in processing rewards, the striatum (Apicella et al., 1991; Schultz, 1998). The strongest projections from both caudal and subgenual portions of the ACC are to the shell of the nucleus accumbens (Yeterian and Pandya, 1991; Lyndbalta and Haber, 1994; Haber et al., 1995). Connections are also evident from all portions of the ACC to the basal amygdaloid nuclei (Amaral and Price, 1984), an area which is implicated in the processing of appetitive responses, and also in the perception of emotion (Adolphs, 2002). Thus, broadly speaking the ACC is connected to areas of association cortex, areas that process rewards and also to areas involved in the processing of emotions.

In the MCC there is some overlap between the areas to which the sulcal area 24c' are connected and the areas to which the gyral areas 24a' and 24b' are connected. These connections also somewhat overlap with the connections of the ACC. Both the sulcal and gyral portions of the MCC are reciprocally connected to areas 9 and 46, much like in the ACC (Pandya et al., 1981; Vogt and Pandya, 1987; Petrides and Pandya, 1999, 2006). Connections are also evident from both the sulcal (24c'/32') and the gyral (24a' and 24b') areas to portions of the medial (areas 10 and 11) and the more lateral (areas 12 and 13) regions of the OFC (Morecraft et al., 1992; Morecraft and Van Hoesen, 1998). However, there are distinctions between the connections of the ACC and the MCC. Notably, the MCC lacks connections to the basal amygdaloid nuclei and to areas TS2 and TS3, areas which are strongly connected in the ACC. Both the sulcal and gyral portions of the MCC project to the core of the nucleus accumbens in the striatum, with very weak connections to the portions of the shell that receives projections from the ACC (Kunishio and Haber, 1994; Haber et al., 1995). In addition, the Ventral Tegmental Area (VTA), an area which is important for processing primary reinforcers (Hollerman and Schultz, 1998; Schultz, 1998; Williams and Goldman-Rakic, 1998), also projects strongly to large portions of area 24c', as well as weakly to areas 24a' and 24b' (Williams and Goldman-Rakic, 1998). However, these connections have only been shown to the more anterior portions of the MCC and unfortunately, it has not been investigated whether connections are also found to more caudal portions of the MCC. Both the gyral and sulcal regions of the MCC are connected to portions of the association cortex in the parietal

lobe, predominantly within medial area 7, but also portions in the intraparietal sulcus (Pandya et al., 1981; Vogt and Pandya, 1987). Thus, there is an overlap between the connections of sulcal and gyral portions of the MCC to areas of association cortex and also to areas that process rewards.

Despite the fact that both sulcal MCC and the gyral MCC are connected to areas of the brain that process rewards, there are also important differences in the areas to which they connect, which suggest that they may have different functional properties. Areas 24a' and 24b' have connections to a portion of the medial wall at the borders of areas 8, 9 and 32', often referred to as paracingulate cortex (Pandya et al., 1981; Vogt and Pandya, 1987; Petrides and Pandya, 2007). The gyral surface of the MCC also has strong reciprocal connections with a region in the depths of the posterior superior temporal sulcus (pSTS), in a region that borders the tempero-parietal junction and also to portions of the most anterior regions of the superior temporal gyrus (Markowitsch et al., 1985; Seltzer and Pandya, 1989; Barbas et al., 1999). The intriguing aspect of these connections is that these areas are considered to be the core-circuit that is engaged when processing others' mental states. There is no evidence of connections from all three of these regions to the sulcal MCC. In addition, the ACC also does not have connections to the same portions of the pSTS or the paracingulate cortex, although there are some weak connections to the similar areas within the temporal poles (Markowitsch et al., 1985; Seltzer and Pandya, 1989). Thus, it would seem that the gyral surface of the MCC is the only portion of the cingulate cortex that has access to information related to the mental states of others.

The sulcus of the cingulate cortex contains three agranular areas, which are referred to as the Cingulate Motor Areas (CMAs). The CMAs consists of the rostral (CMAr), ventral (CMAv) and dorsal (CMAd) regions that are defined as motor areas by their direct projections to the spinal cord (Hutchins et al., 1988; He et al., 1995; Dum and Strick, 1996). In addition, electrical stimulation of neurons in each of these areas results in limb movements (Luppino et al., 1991). The ventral and dorsal CMAs lie in the most anterior regions of the PCC (area 23c). The CMAr lies in the lower bank of the sulcus in the MCC in area 24c'. The CMAs each have strong connections to primary motor, premotor, supplementary motor (SMA) and pre-supplementary motor (pre-SMA) cortices (Showers, 1959; Picard and Strick, 1996; Wang et al., 2001). Outside of the CMAr in area 24c', there are no neurons that project to the spinal cord. However, projections to the primary motor cortex and premotor cortex are found rostrally along the ventral bank of the sulcus (Wang et al., 2001), although such connections are not found along the entire extent of area 24c'. Projections to the SMA and the Pre-SMA are found across the whole length of the sulcal MCC (Wang et al., 2001). Thus, a large proportion of area 24c' has

connections to areas within the motor system. Crucially, no portion of the gyral surface of the MCC has connections to the spinal cord, primary motor cortex or to premotor cortex and very few cells are found that project to the SMA or pre-SMA. Therefore, whilst area 24c' has connections to areas that process rewards like areas 24a' and 24b', its connections strongly implicate the region as being part of the motor system.

Whilst there is segregation between the long-range connections of the different portions of the cingulate cortex, there is considerable evidence that each of the different areas within the cingulate exchange information with each other. Some of the strongest projections appear to be from the gyral MCC to the CMAr in the sulcus (Pandya et al., 1981; Vogt and Pandya, 1987; Morecraft and Van Hoesen, 1998), suggesting that the information processed in the gyrus may have an indirect link to the motor system. Both the CMAr and the gyral MCC are also connected to the most rostral portions of the sulcal MCC and also areas 24a and 24b in the ACC (Pandya et al., 1981; Vogt and Pandya, 1987; Carmichael and Price, 1995). The dominant projections from the subgenual portions of the ACC are to also to areas 24a and 24b, although some weaker projections to posterior portions of the MCC are also found (Pandya et al., 1981; Vogt and Pandya, 1987; Carmichael and Price, 1995; Petrides and Pandya, 2007). This would suggest that information can be exchanged locally within the cingulate cortex, particularly between the gyral MCC and the CMAr.

Broadly speaking, the evidence from tracer studies suggests that there is a similar pattern evident in the connectional anatomy of the ACC, to that which is found in the cytoarchitecture. There are changes in connectivity along the rostro-caudal axis and also distinctions between the connectivity profiles of the sulcus and the gyral surface. In general, anatomical evidence implicates the regions within the ACC in reward-related and affective processes. The MCC also exchanges information with areas that process rewards. However, the gyral surface of the MCC is the only portion of the cingulate cortex that has strong connections to each of the nodes in the core-circuit that is engaged when mentalizing.

1.3.2.2 Human Studies

So far I have outlined a case that supports the notion that the sulcal MCC processes information related to rewards and actions, and the gyral MCC processes socially relevant information. However, this evidence all comes from nonhuman primate research. Until

recently, it was difficult to acquire data that was informative as to the connectional properties of the human brain. However, modern neuroimaging techniques now allow for data that examines anatomical connectivity to be acquired. Techniques such as Diffusion Tensor Imaging (DTI) allow for the claims which have been made based on nonhuman primate tracer studies to be examined in the human brain. DTI is an MRI based technique that enables the measurement of diffusion of water. When water is contained, its direction of diffusion is restricted. When such diffusion occurs in white matter tracts, its diffusion will be restricted in the direction of the pathway (Le Bihan et al., 2001; Le Bihan, 2003). By examining the direction of diffusion of adjacent voxels one can examine whether there are white matter pathways from one brain area to another. Such an approach can be used to examine white matter tracts in vivo, in humans. However, at this point it should be noted that DTI is prone to false negative results (i.e. a real tract is present but not identified) due to the insensitivity of the method in areas where white matter fibres cross (Ramnani et al., 2004a). In addition, although DTI examines white matter pathways in the brain, it does not reveal whether connections exist between neurons at the synaptic level, only whether tracts between areas can be traced.

An alternative method that can be used to examine connectivity between brain areas in humans, is to examine resting-state fMRI data (Beckmann and Smith, 2004; Beckmann et al., 2005). By examining the statistical relationship between BOLD responses in different areas, one can examine whether activity in two areas fluctuates in a similar manner, implicating the two regions as being functionally connected. Obviously, much like DTI, this method is not sensitive as to whether there are connections between areas at the synaptic level. However, it is still informative as to whether the timecourse of activity in two areas is statistically related, and therefore indicative of functional connectivity between areas.

The results of DTI and resting state fMRI studies identify similar patterns of connectivity in the cingulate cortex as those identified in animal tracer studies. Studies highlight a similar change in connectivity extending along the rostro-caudal axis of the cingulate cortex and also suggest distinctions in connectivity between the sulcus and the gyrus (Margulies et al., 2007; Beckmann et al., 2009). Thus, there is some correspondence between human and animal studies in terms of how the cingulate cortex can be subdivided based on its connectivity. In addition, the studies in humans identify similar connectivity profiles of each zone within the cingulate cortex, as those identified in animal studies. However, it should be noted that they do not necessarily highlight the identical connectivity profiles as those identified by tracer studies. Also, there are discrepancies in the connections identified by DTI and by resting-state methods. However, the discrepancies are likely to be due to significant differences in the

nature of the data examined by each method, and the lack of sensitivity of both methods for identifying monosynaptic pathways in the brain. Margulies and colleagues (Margulies et al., 2007) examined resting-state activity in 16 different seed voxels, located across the MCC and the ACC. They found a clear distinction between the areas in which activity covaried with activity within the ACC seed voxels, from those in which activity covaried with activity in the MCC seed voxels. This evidence supports the assertion that the connectivity profiles of the MCC and ACC regions are distinct. DTI methods also support this claim. As can be seen in figure 1.3, Beckmann et al., (2009) performed a DTI-based parcellation of the cingulate cortex, revealing clearly distinct zones in the territory of the ACC and the MCC.

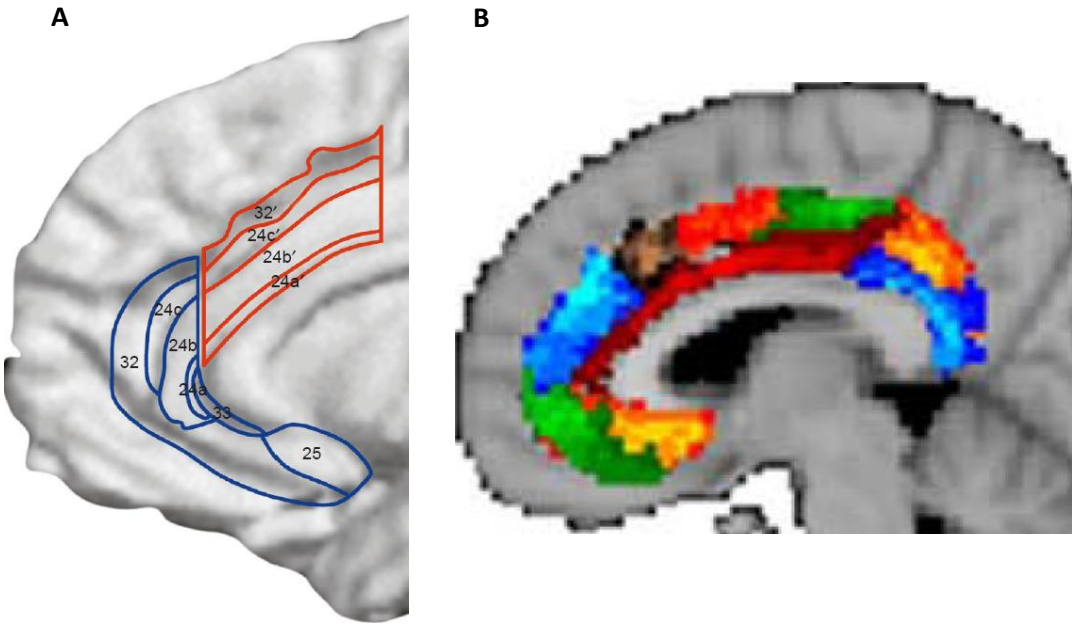

Fig.1.3. (A) Illustration of the cytoarchtectonic zones of the ACC (in blue) and the MCC (in red) taken from Bush et al., (2000), (B) DTI based parcellation of the cingulate cortex by Beckmann et al., (2009). The region immediately adjacent to the corpus callosum in dark red lies on the gyral surface approximating to the location of areas 24a' and 24b'.

Margulies et al. (2007) also found that activity in the MCC covaried with activity in a number of brain areas that were shown to be anatomically connected in the monkey MCC. Activity in the more caudal parts of what they refer to as superior cingulate cortex (putatively lying in the sulcus of the MCC) was statistically related to activity in areas that process rewarding stimuli, including the OFC, the striatum, the PCC and the frontal pole. Activity was also statistically related to the premotor and primary cortices. Beckmann et al. (2009) also identified tracts between a dorsal cingulate region (extending into the sulcus, putatively in MCC) and areas that contain reward sensitive neurons, including parietal cortex and the striatum, as well as connections to the primary motor and premotor cortices. Each of these areas have been found

to be connected to area 24c' in nonhuman primates. These connections support the notion that the sulcal MCC is important for processing both rewards and actions.

The results revealed by methods investigating the connectivity of the MCC gyrus in humans is less clear. A DTI study by Beckmann et al. (2009) found that although an area corresponding to the gyrus of the MCC showed a distinct pattern of connectivity to the rest of the cingulate cortex, it did not show strong connections to any specific area. However, in their examination of resting state data, Margulies et al. (2007) reported that activity in the "inferior cingulate cortex" which putatively corresponds to the gyral MCC, statistically varied with activity in reward processing areas including the striatum, the OFC and portions of the parietal cortex. Importantly this study also reported activity in the paracingulate cortex and the posterior superior temporal cortex that covaried with activity in the gyral MCC. As stated in section 1.2, these two areas constitute part of a core-circuit which is engaged when processing others' mental states. Thus the functional connectivity of the gyral MCC in humans does show some correspondence with the regions that are also anatomically connected in animals. Therefore, bearing in mind the sensitivity of these methods, the evidence they provides supports assertions that were made above; The MCC has separate gyral and sulcal regions which both process reward-related information. The connections of the sulcus suggest that it processes information which guides actions, whereas the connections of the gyrus suggest that it processes social information. Later in this chapter, I will argue that the gyral surface and the sulcus in the MCC process similar information during reward-guided decision-making. However, the gyrus processes information about the reward-guided decisions of others.

1.4 The ACC and Social Cognition.

As discussed in section 1.2, the evidence from patients with lesions that included portions of the ACC, suggests that the ACC is involved in processing social information. However, the damage in such patients is not informative as to what specific cognitive processes are disrupted and where such processes are localized within the cingulate cortex. In this section, I present evidence from experimental lesion studies that show that a specific sub-region of the ACC is involved in social behaviour and in processing social stimuli (Rudebeck et al., 2006a). These studies show parallels with the anatomical evidence reported in the previous section, highlighting the gyral surface, rather than the sulcus, as a candidate for processing social information. It should be noted that in these studies, lesions are made to both the sulcal and gyral portions of the ACC and the MCC, but they are referred to as the ACCs, (cingulate sulcus) and the ACCg (the cingulate gyrus). I will use this nomenclature throughout this section, in order to remain consistent with the authors' discussion of their results. I will, however, discuss anatomical locations when they are relevant.

The connections of the ACCg to the core-circuit implicated in theory of mind processing raise the possibility that the ACCg may be similarly engaged when processing others' mental states. In the second part of this section I will discuss this possibility. However, as will become clear, there is an absence of any functional imaging evidence implicating the region in processing others' mental states. In the final section, I will attempt to reconcile the contrasting results from animal and human studies, by reporting evidence that the ACCg may process information about others in the same manner that the ACCs processes information about oneself.

1.4.1 Animal Studies

Over the last decade a number of paradigms have been developed that are able to measure the extent to which an animal is evaluating and processing information about another. In a seminal study by Hadland et al. (2003), the social behaviours of three macaque monkeys were recorded following lesions to the ACC and compared with three macaque monkeys without such a lesion. The location of the lesions extended across the ACCg and the ACCs, in both MCC and ACC regions, and in one case extended across a portion of the paracingulate cortex. However, in all three cases a portion of the gyral surface in the MCC was ablated. The behaviour of the monkeys was examined whilst housed in a cage with two other monkeys

without an ACC lesion. The ACC lesioned monkeys spent less time stationary next to another monkey, less time interacting with other monkeys and performed fewer vocalisations directed at other monkeys than the control group. In a second experiment, lesions were made to the ACC in the monkeys that had acted as control monkeys in the first experiment. The monkeys that were lesioned in this second experiment showed the same pattern of behaviour as the group who had lesions in the first experiment. Importantly, these behavioural changes could not be the result of general akinesia as the animals spent a significantly increased amount of time playing with a toy object following the lesion (Hadland et al., 2003). This study supported the tentative conclusions that can be drawn from the patients with acquired sociopathy discussed earlier in this chapter. That is, a portion of the cingulate cortex may be important for processing information that guides social behaviours. Similar findings have been found in rats. Lesions to the entire ACC in the rat have been shown to reduce the amount of time that the animals spend engaging in behaviours which constitute social interactions, such as sniffing other rats (Rudebeck et al., 2007). This would therefore suggest that the ACC may be important for guiding social behaviour in several different species.

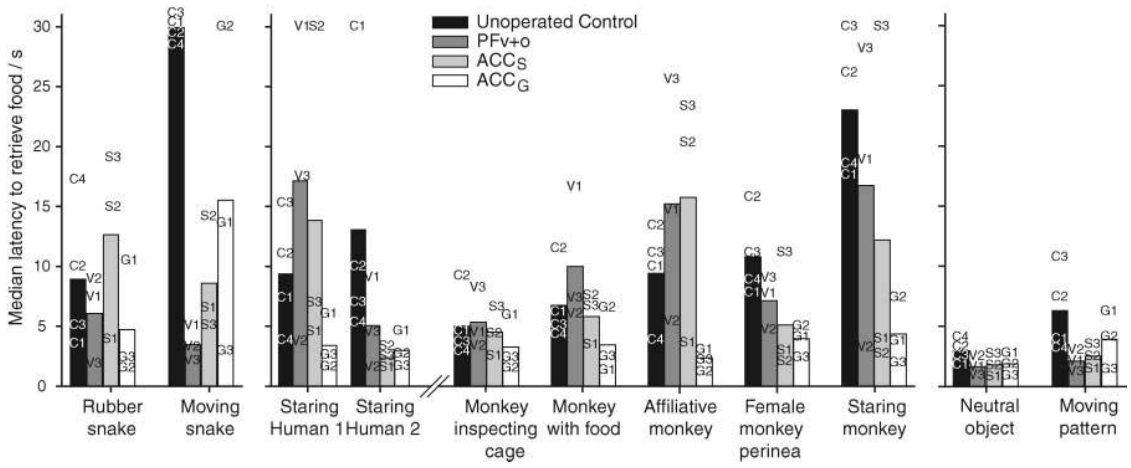

Fig.1.4. Figure taken from Rudebeck et al.,(2006). The graph shows the latencies before reaching for a food item, when the food is presented with another stimulus. Lesions to the ACCg resulted in significantly reduced latencies when the food was presented at the same time as a social stimulus. Lesions to the ACCs and the orbitofrontal cortex (PFv+o) did not show such significant reductions in the latency.

Rudebeck and colleagues (Rudebeck et al., 2006a) performed a set of experiments that extended the results of Hadland et al., (2003). In their study, they examined the latency that occurs before reaching for a food item when the food was presented simultaneously with

different types of social and emotional stimuli. Monkeys without lesions, with lesions to the lateral OFC or lesions to the ACCs, showed a similar latency before reaching for a food item when it is presented with a social stimulus. In contrast, monkeys with a lesion specifically to the ACCg (including the entire gyrus in the ACC and a portion of the gyrus in the MCC) showed significantly decreased latencies, in the presence of multiple different types of social stimuli. In fact, the latencies in these monkeys were the same as those exhibited by the control monkeys when the food was presented at the same time as a neutral stimulus (see fig.1.4). In addition, the ACCg lesions also resulted in a reduction in communicative behaviours such as lip smacking and communicative vocalizations, which were not reduced with OFC or ACCs lesions. This study therefore supports the notion, that it is the ACCg that processes social information and guides behaviours during social interactions. This study therefore corroborates with the anatomical evidence that the ACCg plays an important role in social behaviour.

1.4.2 Does the ACCg process others' mental states?

As outlined in section 1.3, the ACCg has strong reciprocal connections to the paracingulate cortex, the pSTS and the temporal poles. These areas are considered a core-network that processes others' mental states. This suggests that the ACCg has access to information about others' mental states, which in turn implicates this area as an additional candidate for processing the mental states of others. One might therefore expect that tasks in which subjects are processing others' mental states, which activate the paracingulate cortex and the pSTS, might also activate the ACCg. However, the majority of neuroimaging studies that have investigated Theory of Mind processing do not show any mentalizing related responses in the ACCg (Fletcher et al., 1995; Vogeley et al., 2001; Berthoz et al., 2002; Berthoz et al., 2006; Saxe et al., 2006; Young et al., 2007; Hooker et al., 2008; Aichhorn et al., 2009; Apperly and Butterfill, 2009; Young et al., 2010a; Zaitchik et al., 2010; Cloutier et al., 2011; Rothmayr et al., 2011; Schnell et al., 2011). There are only a limited number of studies that have reported activity in any region of the ACC during conditions in which subjects are attributing mental states to others and these either fall in the most anterior portions of the ACC or in the sulcus (Brunet et al., 2000; Chaminade et al., 2007; Sommer et al., 2007). As such, these isolated studies do not report activity in the portion of the ACCg (in the gyral MCC) that in this thesis is argued to process social information.

The absence of many studies reporting activity in the ACCg when processing others' mental states is not restricted to studies that specifically examine Theory of Mind abilities. A recent meta-analysis of over 200 studies which investigated the processing of social information in many different contexts, also failed to find many studies which reported activity in the ACC. This meta analysis included studies which investigated the processing of others' negative and positive emotions, social norms and others' personality traits (Van Overwalle and Baetens, 2009). Within and across these different categories, the gyral surface of the MCC was not implicated in processing social information.

There is one notable exception to the studies that do not find ACCg activity when mentalizing. In the study by Walter et al., (Walter et al., 2004) that has already been reported in section 1.2, activity was found in the paracingulate cortex when subjects processed the communicative intent of others. In their study, they reported activity in the ACCg during the control condition, where the mental states of others' could be processed but their mental states related to actions that were not related to social interactions. For example, the ACCg was activated when subjects observed another about to change a light bulb. Thus, the subjects could infer the mental states of another at the time that the predictions about others' actions were being processed. In later sections of this chapter I claim that the function of the ACCg may be to process others' predictions about their actions, which may therefore account why this study identified such activity.

In summary, despite strong connections to areas of the brain which are engaged when processing another's mental state, the majority of neuroimaging studies in humans which investigate mentalizing abilities fail to find any ACCg activity occurring when attributing mental states. This contrasts with the evidence from lesions in both humans and also in animals which have implicated the ACC and particularly its gyral surface as important for social behaviour.

1.4.3 Reconciling Human neuroimaging and Animal lesion studies

Starkly contrasting conclusions can be drawn from the neuroimaging literature investigating mentalizing processes and the lesion studies investigating disruptions in social behaviour. Can the two sets of evidence be reconciled? An obvious way to reconcile these two conflicting strands of evidence would be to take single-cell recordings from the ACCg when monkeys are interacting with others. However, to the best of my knowledge, there have been no studies which have recorded from the ACCg during social interactions. In fact, there is absence of studies that have documented the functional properties of neurons on the gyral surface of the MCC at all.

Another possible reconciliation, given that differences occur between species, could come from an evolutionary perspective. In section 1.3 I highlighted that the ACCg had evolved in humans, with a significantly increased density and quantity of spindle neurons in the ACCg in comparison to any other species. Thus, it could be argued that as a result of evolutionary pressures the ACCg is no longer engaged by social information. However, this argument lacks any empirical support and, in fact, contradicts the only evidence relating to the evolution of the ACCg. Indeed, spindle neurons (see section 1.3.1) are found only in species that engage in complex social interactions and the complexity of social interaction correlates with the quantity of this type of neuron (Nimchinsky et al., 1999; Butti et al., 2009; Allman et al., 2011). Thus, although this is only a correlation, the only evidence examining the evolution of the ACCg argues for it having a more prominent role in processing social information in humans.

An important fMRI study (for the field in general, but particularly for the work in this thesis) by Behrens et al. (2008) offered an alternative reconciliation. They noted that the ACCg and the ACCs, despite different connectivity profiles, share connections to systems that are engaged in processing the rewarding outcomes of decisions. They hypothesised that the ACCs and the ACCg might compute similar information about the outcomes of one's own and others' decisions respectively. Subjects performed a two-choice decision-making task which required them to track the probability of receiving a reward and the stochasticity of the ratios between two different rewarding options, i.e. track the volatility of the delivery of rewarding outcomes. During the experiment the reward schedules changed, such that in some periods the reward schedules were more stochastic than in others. On each trial, they also received advice from another participant (a confederate) about which was the better of the two options. They tested the hypothesis that the ACCs would code for how volatile the association was between

their chosen options and the rewarding outcomes and the ACCg would code for how volatile the association was between the confederate's chosen options (the advice) and the rewarding outcomes. They reported activity in the ACC at the time of the outcomes. Activity in the ACCs and the ACCg was found to vary with degree to which people should learn from their choices and the confederate's choices respectively, as a function of the volatility of the feedback (see fig.1.5). This study suggests that the ACCs and the ACCg process the same information about one's own and other's decisions respectively.

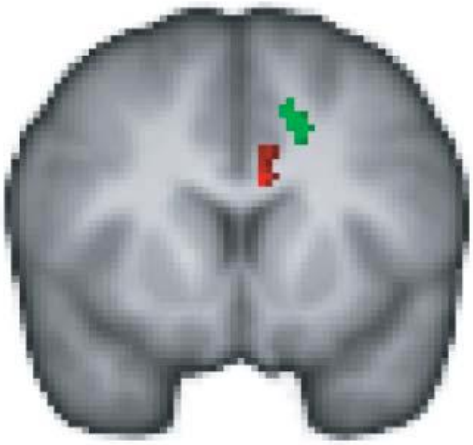

Fig.1.5. Figure taken from Behrens et al.,(2008. Activity shown in the ACCs (green) and the ACCg (red). Activity in these two areas was found at the time of the outcome of decisions. Activity in the ACCs coded for the volatility of decision-outcome contingencies. Activity in the ACCg coded for the volatility of the advice provided by a confederate. Both clusters were putatively in the MCC.

Importantly, the nature of the activity that was identified in this study may also be informative as to the absence of studies which report this area as being engaged when mentalizing. In this study, the ACCg was shown to process the level of volatility of another's advice. Activity in this area varied with a computation in a learning model (computational models will be discussed in section 1.6 in more detail). Thus, the response in the ACCg was parametric and occurred time-locked to a specific event during decision-making. As discussed earlier, studies investigating theory of mind processing typically use block designs, where activity is aggregated across a whole vignette. However, if the ACCg processes social information only at the time of particular events and activity in this area varies parametrically, subtraction designs between activity aggregated over blocks would not have enough power to identify activity in this area.

The study by Behrens et al., (2008) therefore provides important evidence to suggest that the ACCg processes information about others during social interactions. A handful of other studies have shown activity in the ACCg when subjects are interacting with others. In a similar manner, they show that the ACCg is engaged when subjects' monitor the decision-making of others and that activity in this area occurs time-locked to specific events. Such studies have shown that

the ACCg in the brain of a subject is engaged when processing cues that instruct others to make a decision (Tomlin et al., 2006; Baumgartner et al., 2009), when the response another has made is observed (King-Casas et al., 2005; Baumgartner et al., 2009) and when the outcomes of another's decision is monitored (Shane et al., 2008; Apps et al., accepted). The results of human studies and animal studies may therefore be reconciled by the notion that the ACCg is engaged specifically when monitoring the decision-making of others. In this thesis, four studies will be presented that examine what specific contribution the ACCg makes to understanding the decision-making of others.

1.5 Theoretical Accounts of ACC Function

The absence of a substantial body of literature implicating the ACCg in processing social information has also resulted in an absence of a theoretical account of its contribution to social cognition. Behrens et al., (2008) suggested that the ACCg may process similar information about others, which the ACCs processes about oneself. If analogous information is processed in these two areas, the ACCg may conform to the same theoretical principles as the ACCs. In contrast to the absence of theoretical accounts of the role of the ACCg in social cognition, there have been several theoretical accounts of the function of the ACC as a whole. In this section, I will outline several theories of ACC function and review the findings of studies that have used electrical stimulation, single-unit recordings, lesions and functional imaging to test them. Once the functional properties of the ACCs have been outlined, it will be possible to hypothesise what contribution the ACCg makes to social cognition. In this section I will continue to use the terms ACCg, ACCs and ACC to refer to both the MCC and ACC, in order to remain consistent with the authors' discussion of their results.

1.5.1 Historical accounts of ACC function

For the majority of the twentieth century there were two opposing and seemingly parallel viewpoints on the function of the ACC. One viewpoint was that that the ACC is involved in highly abstract processes (Devinsky et al., 1995). This line of thinking arose due to changes in behaviour that occur when the ACC is damaged. Patients who have undergone a cingulotomy often report an absence in motivation and spontaneous response in the months following surgery, sometimes lasting up to 12 months (Cohen et al., 2001). Such a behavioural change is also evident and even more striking in patients who suffer from akinetic mutism, which can occur following ACC damage (Buge et al., 1975; Nemeth et al., 1988; Devinsky et al., 1995). Patients' who suffer from this debilitating disorder, have the ability to speak and respond to others, as well as the ability to make normal movements. However, they do not perform any action or speech spontaneously without external input (Buge et al., 1975; Nemeth et al., 1988). Patient E.V.R suffered from akinetic mutism following the removal of a large meningioma on the medial wall of the frontal lobe that was removed at age 8. Following the removal of the tumour there was considerable damage affecting the left hemisphere, including a large portion of the ACC. Although E.V.R suffered from many deficits for a considerable period after the surgery, she recovered from the akinetic mutism and was later able to offer considerable

insight into the syndrome. The patient reported that she did not perform actions or speech because she felt no "will" to respond, as if her actions had no value and "nothing mattered". Some suggested on the basis of this evidence, that the ACC must be a region that is important for conscious experience (Damasio and Van Hoesen, 1983; Damasio and Geschwind, 1984). However, it is important to exercise caution when trying to use such evidence to ascribe function to an anatomical region. Lesions that are the result of brain damage are not specific to any one gross anatomical or cytoarchitectonic region. In addition, lesions may result in diaschisis, disrupting processing across a distributed network. Thus strong conclusions should not be drawn from such data.

In contrast to the viewpoint that the ACC is involved in higher-order processes, there is considerable evidence that the ACC is engaged in very low-level visceral processes. Studies in animals that have stimulated the cingulate gyrus report changes in heart rate, blood pressure, and aggressive behaviours (Kaada et al., 1949; Terreberry and Neafsey, 1983; Hurleygius and Neafsey, 1986; Neafsey et al., 1986; Terreberry and Neafsey, 1987; Frysztak and Neafsey, 1994). Neuroimaging studies have also shown that activity in the ACC responds to breathing rate and also levels of hunger (Brannan et al., 2001; Liotti et al., 2001). Each of these processes is visceral and contrasts with the neuropsychological evidence above.

So, how can the ACC be a structure that processes low-level, visceral, autonomic processes and also the most complex, abstract information necessary for human conscious experience? Some have attempted to reconcile these two accounts, suggesting that the ACC is engaged in the planning, execution and generation of behaviours (Paus, 2001). Some have also suggested that the ACC is involved in translating high-level information to low-level systems or vice versa (e.g. transmitting emotional responses into autonomic systems or translating emotions into cognitive systems) (Devinsky et al., 1995; Morecraft and Van Hoesen, 1998; Bush et al., 2000). Yet, such descriptions are highly unspecific, unfalsifiable and possibly could be used as a description of many areas of the brain. As stated in section 1.3, part of this confusion arises from the use of the term ACC when referring to an anatomically, and therefore functionally, heterogonous area. Important clues as to the functional properties of the two areas that are examined in this thesis (the sulcus and the gyrus of the MCC) can be found in neuroimaging studies in healthy human adults. The findings of such studies may reveal important clues as to how the ACCs may contribute to behaviour and therefore be informative as to how the ACCg might contribute to social cognition.

1.5.2 Cognitive and Affective divisions in the ACC

Whilst historical accounts of ACC function have used evidence from lesions and electrophysiology, more recent theories have been developed around the findings of neuroimaging studies in humans. Functional imaging research, has implicated the ACC in processing a broad range of different information, including: pain (Bush et al., 2000), the detection of response errors (Carter et al., 1998), conflict monitoring (Botvinick et al., 1999; Kerns et al., 2004) and action selection (Rushworth et al., 2004). However, neuroimaging research has revealed that in the cingulate cortex there are functional divisions that correspond with the anatomical divisions. Specifically, there are changes in function along the rostral-caudal axis. Broadly speaking, the anatomically defined ACC responds during the processing of emotions, whereas the MCC responds during cognitive processes.

One convincing body of evidence that highlighted an ACC-MCC division in the processing of affective and cognitive information was provided from studies using the Stroop test (Stroop, 1935; Bush et al., 2000). The Stroop task is a long-standing psychological test of the ability to process incongruent stimuli. Subjects are required to read aloud a list of colours, that are either congruent or incongruent with the colour in which the text is printed (i.e. red and blue are congruent, but red and blue are incongruent). Behaviourally, an increase in reaction times occurs when a words' meaning is incongruent with the font colour. In addition, on some trials subjects make an error, reporting the colour of the text rather than the written word. This task has also been adapted to examine the processing of incongruent emotional and cognitive stimuli. In the emotional counting Stroop task, subjects are required to count the number of words on the screen and then identify when a word is emotionally different from the other words. Another adaptation is the cognitive counting Stroop task, where subjects count the number of numbers that have been presented in a series and have to detect when the written number is incongruent from the number of stimuli which they have counted (Davis et al., 2005). These tasks can be used to examine the processing of emotional stimuli and cognitive abilities.

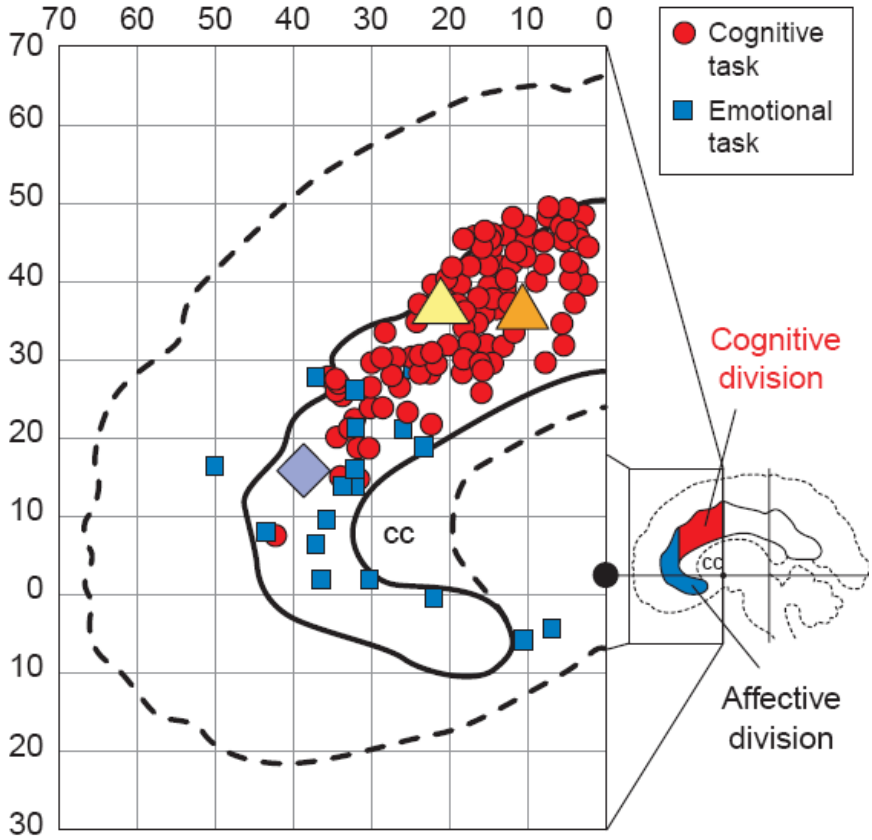

Fig.1.6. Figure taken from Bush et al., (2000). A meta-analysis of studies investigating the processing of emotion (blue squares) or cognitive processes (red squares). Emotion processing studies activated the caudal ACC, whereas cognitive studies activated the MCC.

Two neuroimaging studies that were conducted on the same nine subjects, found that activity during the emotional and cognitive Stroop tasks, activated different portions of the ACC (Bush et al., 2006; Whalen et al., 2006). When the emotional Stroop task was performed, a region was activated in the caudal regions of the anatomically labelled ACC (i.e. in areas 24a-c). In contrast, in the same subjects, a more dorsal region in the MCC was activated when subjects performed the cognitive Stroop (i.e. in areas 24a'-c'). In addition to the evidence from these two studies, a meta analysis of neuroimaging studies investigating cognitive and emotional processes supports the assertion that there is correspondence between MCC-ACC and cognitive-affective cingulate divisions (Bush et al., 2000). As can be seen in figure 1.6, the peak coordinates of studies investigating cognitive processing are almost all in the MCC, whereas the affective tasks peak coordinates predominantly lie rostral to the MCC in the ACC.

The meta-analysis conducted by Bush et al. (2000) revealed that there is an important distinction between information processing in the ACC and the MCC, with the MCC being engaged by cognitive processes. Intriguingly, the majority of the reported activations in tasks

that investigate cognitive processing lie superior to the gyrus, in the sulcus of the MCC. This corroborates with the anatomical evidence presented earlier, which highlighted the rostro-caudal, sulcal-gyral distinctions in morphology and connectivity. In previous sections I have suggested that the gyrus in the MCC processes social information. Is the ACCg engaged when monitoring others' cognitive processes? This hypothesis was tested in an fMRI experiment by Shane et al. (2008). Subjects were required to perform, or observe a third-person performing, a go/no-go task. Go/ no-go tasks are considered a cognitive task as they require the ability to cancel and inhibit a prepotent response or to perform an unplanned response. Shane et al., (2008) reported activity in the ACC, in a similar portion of the sulcus in the MCC area engaged during the cognitive Stroop task, whenever an erroneous response or absence of a response occurred, regardless of whether the trial was performed by the subject or performed by a third-person. However, an additional portion of ACCg, on the gyral surface of the MCC, responded exclusively to the erroneous inhibitions and erroneous responses of the third-person (Shane et al., 2008). Thus, when the cognitive processes of another are erroneous, activity is found in the ACCg. These studies offer tentative support for the view that the ACCg is engaged when subjects are monitoring the behaviour of others, that reveals important information about their cognitive processing. Thus, the MCC can be considered as an area engaged during cognitive processing, with the gyrus sensitive specifically to others' cognitive processes.

1.5.3 Theoretical accounts of 'Cognitive' ACC?

As stated in previous sections, there is an absence of a theoretical account of how the ACCg contributes to social cognition. I have argued that the ACCg is engaged by similar cognitive processes as the ACCs, but processes this information when monitoring the behaviour of others. Thus, in order to understand how the ACCg might contribute to social cognition, it is first important to understand the cognitive processes that are performed in the ACCs

As the aim of this thesis is to understand how the cingulate cortex contributes to social cognition, and not to the processing of emotions, future discussion in this chapter will be restricted to discussions about the functional properties of the MCC. The research I have reported thus far, has not examined what cognitive processes engage the MCC. In this section, I will outline three theoretical accounts of cognitive ACC function. In later sections, I will argue that these theories cannot account for all of the functions that can be ascribed to the cognitive

ACC division. However, these theories have been pervasive in the literature and research continues to test their underlying assumptions. Additionally, the research investigating these theories has revealed important information about the functional properties of the cognitive ACC division that any framework must incorporate. For these reasons a discussion of these accounts is pertinent. It should be noted that alternative theories have been proposed for ACC function. For instance, some have suggested that the ACC is involved in consciousness (Devinsky et al., 1995) and others have suggested that the ACC is involved in attention (Corbetta et al., 1991). However, considerable research has highlighted their insufficiency and so they will not be discussed here.

1.5.3.1 Conflict Monitoring

The most prominent and comprehensive theory of ACC function to date, suggests that this area is crucial for detecting conflicts in information processing (Botvinick et al., 1999; Barch et al., 2000; Botvinick, 2007). Conflict monitoring stems from an idea from cognitive psychology, that novel or unexpected items require prepotent responses to be overridden. Many have suggested that the ACC is well placed to signal such conflicts as it has such widespread connections to areas of association cortex. It has therefore been proposed that the ACC signals when there is conflict between two different streams of information processing and also, that the greater the level of conflict, the greater the response in the ACC will be (Botvinick, 2007). This theory has considerable support, particularly from studies using neuroimaging methods in humans.

Evidence for a role of the ACC in conflict monitoring was first presented by studies using the Stroop task that has already been described (Bush et al., 2000; Kerns et al., 2004; Bush et al., 2006; Whalen et al., 2006). In this task, when a novel or unexpected item is presented ("blue") following the word that has been repeated ("blue") , there is conflict between the prepotent response ("blue") and the response that is now required ("red"). Thus, it is suggested that at the time of incongruent stimuli, there is conflict between previously repeated responses to the congruent stimuli and the new response which must override the prepotent one. In line with this suggestion, activity at the time of an incongruent stimulus is greater in dorsal ACC regions (putatively in the sulcus in the MCC) than at the time of a congruent stimulus in the Stroop task (Kerns et al., 2004; Bush et al., 2006; Nee et al., 2011). Other neuroimaging studies which have used very different tasks, which also result in conflicts in information processing, have

reported activity in the MCC. Experiments using the Go/No-Go task, where a subject is required to make a response on many trials ("Go" trials) before a trial occurs where they are required not to make a response ("No-Go"), have reported activity in the ACC. In particular, activity is greater in the ACC on the trials where the unexpected event occurs (Kawashima et al., 1996; Menon et al., 2001; Rubia et al., 2001). This has been interpreted as indicative of conflict occurring between the prepotent "Go" response and the now required "No-Go" response or vice versa. In fact, many different tasks that require a prepotent response to be overridden by another, report activity in the dorsal ACC when such a conflict is present (Kerns et al., 2004; Sohn et al., 2007; Fan et al., 2008; Pochon et al., 2008; Nee et al., 2011). Activity in the ACC also increases with the number of possible actions that can be taken, which is interpreted as being indicative of the increased level of conflict between the different possible actions. Such an effect has been demonstrated in decision-making paradigms, where activity in the ACC increases at the time of instruction cues when the number of alternative possible choices that are available increases (Braver et al., 2001; Fan et al., 2008; Pochon et al., 2008).

Whilst there is a considerable body of neuroimaging research which supports the conflict monitoring account of ACC function, there is little support from studies investigating the functional properties of neurons in this area. Studies in non-human primates have found little evidence of neurons in which activity is sensitive to changes in conflict (Amiez et al., 2005; Sallet et al., 2007; Quilodran et al., 2008; Kennerley et al., 2009; Kennerley and Wallis, 2009b, a; Hayden and Platt, 2010; Hayden et al., 2011b). Indeed, several of the studies that use single-cell recording methods to examine the functional properties of neurons in the ACC find neurons which change their firing rate in response to stimuli, when the level of conflict has been experimentally controlled. However, it is possible that conflict monitoring processes only occur in human ACC, explaining the absence of such processing in non-human primate ACC. Rare single-unit recording data in humans, that recorded from the sulcus in dorsal MCC support this claim, showing that there are neurons which increase their firing rate with level of cognitive conflict during a Stroop task (Davis et al., 2005).

Studies that examine the processing of conflict typically activate the sulcal MCC. Is the ACCg engaged when monitoring others' conflict? In the study by Shane et al. (2008), it was noted that activity in the ACCg occurred on the trials in a Go/No-Go task when subjects observed another failing to correctly inhibit, or correctly perform an action. This could be interpreted as the subject monitoring the conflict of a third-person. As such, this could suggest that the ACCg is engaged when processing others' conflict.

1.5.3.2 Error Detection

Whilst many proponents of the conflict monitoring viewpoint suggest that the activity on incongruent trials during a Stroop task is a signature of conflict, there are alternative interpretations of such activity. One of the most prominent is that an incongruent trial in a Stroop task signals an error in the prediction of an upcoming stimulus (Bush et al., 2000; Holroyd and Coles, 2002). It is suggested that the Stroop task engages the ACC, as this region has a general role in coding for erroneous responses.

The majority of the evidence which supports the error detection perspective comes from EEG studies. The ACC is believed to be the dipole source of an electrophysiological response that occurs whenever an erroneous response is made, known as the Error-Related Negativity (ERN). A large number of studies show that the ERN occurs when an action does not result in the desired goal (Braver et al., 2001; Menon et al., 2001; Nieuwenhuis et al., 2001; Holroyd and Coles, 2002; Holroyd et al., 2003; Frank et al., 2005). fMRI studies have also found a portion of the ACC is engaged when processing erroneous responses. The location of error related activity in fMRI studies appears to closely correspond to the localization of the source of the ERN and typically falls in the ACCs (Carter et al., 1998; Kiehl et al., 2000; Braver et al., 2001; Garavan et al., 2002; Holroyd et al., 2004). Thus, the evidence in humans broadly supports the claim that the ACC is engaged when responses are erroneous.

The study by Shane et al., (2008), could be considered as evidence supporting the notion that others' erroneous responses are processed in the ACC. As already discussed, they reported activity in the ACCg when another's response or absence of a response was erroneous. Proponents of the error detection account would suggest that these trials in a Go/No-Go task reflect erroneous responses, rather than conflicts in information processing, Thus, the error detection account of the results of Shane et al., (2008) would suggest that the ACCg processes others' erroneous responses in the same manner as the ACCs processes one's own.

Evidence in support of the error detection account of ACCs function is also provided from animal neurophysiology data. It has been shown that there are neurons in this area that respond exclusively when an action does not lead to an expected outcome (Matsumoto et al., 2007). In addition, other studies have also shown that there are neurons that respond when a predicted reward is not received (Amiez et al., 2005; Sallet et al., 2007). Neurons which exhibit such a response profile conform closely to the predictions made by the error detection account.

1.5.3.3 Response Selection

In monkeys, the CMAr has direct projections to the spinal cord and also to the premotor and primary motor cortices (Dum and Strick, 1996; Picard and Strick, 1996). Electrical stimulation of neurons in this region, elicits complex multi-joint movements (Luppino et al., 1991). As a result, many have suggested that the most important functional property of the ACCs must be that it is involved in selecting and guiding actions. In humans, although the location is much debated, there is a homologous region which is often called the Rostral Cingulate Zone (RCZ). fMRI studies using the most basic of motor paradigms, such as finger tapping, reliably activate the ACCs (Kawashima et al., 1999; Buchel et al., 2002; Ullen et al., 2003), suggesting that the RCZ may be homologous to the CMAs.

If the ACCs is involved in the generation of actions, one would expect that the lesions to the ACC would inhibit the willingness to perform actions. As has been discussed, this is the case in patients with akinetic mutism, which can occur following damage to the ACC. These patients report an absence in the motivation for performing an action (Damasio and Van Hoesen, 1983). More specifically, lesions to the ACC and also separate lesions to the adjacent Supplementary Motor Areas (SMA) in a monkey, result in a decrease in the number of actions that are performed without an external cue (Thaler et al., 1995). Interestingly, a number of fMRI studies also implicate the SMA (or pre-SMA) and the RCZ in generating actions. Both of these areas have been found to be activated before the onset of an action that occurred without any external cue (Lau et al., 2004b; Lau et al., 2004a; Lau et al., 2006; Passingham et al., 2010). However, whilst the SMA appears to be involved predominantly in self-generated actions, the ACCs is engaged both by internally and externally cued actions (Lau et al., 2006). Thus, there is considerable evidence that highlights the ACCs and particularly the CMAr as important for the selection of an action. However, it is important to note that unlike the error detection and conflict monitoring accounts, there has been no reported evidence that highlights the ACCg as a region which is engaged when processing others' actions.

1.5.3.4. Critique of the three theories of ACC function

In the previous sections, I have presented the cases in support of three of the most prominent accounts of the functional properties of the ACC. In this section I critique these as accounts of the functional properties of the ACCs. Whilst the studies reported are not examining information processing in the ACCg, parallels are being drawn between the functional properties in the ACCs and the ACCg later in the chapter. The discussion in this section is therefore important for understanding the functional properties of both areas.

An important consideration, before critiquing these theories as explanations of the functional properties of the entire ACC, is where the response selection, error detection and conflict monitoring processes are localized in the same portions of the ACC. If these processes can be localized to the same portion of the ACC, then they can be considered as opposing theoretical accounts. Recently, a meta-analysis of fMRI studies by Beckmann et al., (2009) examined the correspondence between the location of several different functions which are known to be processed in the ACC and a DTI based parcellation of the cingulate cortex. This study grouped previous research that reported activity in any portion of the entire cingulate cortex (inclusive of the PCC) into one of 7 categories: emotion, reward, conflict, error detection, motor function, memory and pain. This meta-analysis showed that the peak results of the studies in the error, motor and conflict categories each overlapped with the same cluster from the DTI parcellation. However, there was variability in peak coordinates, with some of the peak coordinates from the studies investigating conflict lying anterior to the error and motor peak results. In addition, some of the peak coordinates from the motor studies lay posterior to the peak results from the motor and error studies. A recent fMRI study examined activity in the ACC of subjects who were performing three different tasks which tested the hypotheses of each of the three theories (Nee et al., 2011). This study examined activity on trials where there was conflict, when subjects were required to switch to an alternative action (task-switching) and when monitoring the outcome of erroneous actions. They reported activity in the ACCs during all three tasks, supporting the notion that the ACC is engaged by all three processes. They also reported considerable overlap in the location of activity for each of these processes. However, the cluster that was activated during the outcome monitoring task lay posterior to the cluster that responded to the conflict and task-switching tasks. This would suggest that there are distinct sub-regions within the ACC that are engaged when selecting actions, monitoring conflict and when detecting errors. However, these results contradict with the results of the meta-analysis by Beckmann et al., (2009), who showed that motor tasks typically

activate slightly more posterior regions of the ACC than error detection or conflict monitoring tasks. Previous meta-analyses have also attempted to delineate zones in the ACC which process each of these types of information (Koski and Paus, 2000; Paus, 2001; Margulies et al., 2007; Nee et al., 2007). These have also failed to find any consistent delineation between these processes.

As there appears to be overlap in the anatomical localization of conflict, error and response selection processes, the theoretical accounts can be considered as opposing explanations of ACC function. However, there is considerable evidence to suggest that none of these theories can account for all of the data that speaks to the function of this area. The first criticism that can be directed at all three theories is that none of the studies can account for all of the evidence from single-unit recording studies. As stated above, there is little evidence to suggest that there are neurons in the ACC that are sensitive to conflict in the monkey. By contrast, there is considerable evidence to suggest that the ACC contains neurons that signal when a response was erroneous (Amiez et al., 2005; Matsumoto et al., 2007). In addition, stimulation of neurons in the CMAr of the ACC results in limb movements, highlighting how this region is important for the generation of actions (Luppino et al., 1991). However, these are not the only types of information that are processed in this same region of the ACC. For example, it has become well documented that in nonhuman primates there are neurons in the same region (in the CMAr) which are sensitive to reward magnitude, reward probability and the amount of effort that is required for a reward (Shidara and Richmond, 2002; Sallet et al., 2007; Quilodran et al., 2008; Kennerley et al., 2009; Kennerley and Wallis, 2009a). These neurons increase their firing rate to such information at the time that actions are instructed and also at the time the outcomes of actions are monitored. It could be argued, that the ACCs has evolved and therefore may process different information in the human than it does in the monkey, accounting for the discrepancy between nonhuman and human data. However, there is evidence to counter this claim, with a number of neuroimaging studies showing that activity in the ACCs varies with the magnitude and probability of rewards in humans (Knutson et al., 2000; Knutson and Cooper, 2005; Rolls et al., 2008; Smith et al., 2009). Therefore, at the time that actions are performed and at the time that outcomes are evaluated, the ACC codes for information related to rewards. None of the three theories outlined in this section can account for the processing of reward-related information at these two points in time. The error detection theory cannot account for the evidence that suggests that the ACC processes reward-related information, before an action has been made and the outcome evaluated. The response selection viewpoint cannot account for the substantial body of literature that shows

that the ACC codes reward-related information at the time of the outcome of an action. Finally, the conflict monitoring viewpoint cannot account for the fact that neurons in the ACC code for the reward associated with an action, when no alternative action can be selected i.e. when there is no conflict.

In summary, these three prominent accounts of ACC function are all supported by evidence, indicating that the ACCs may have an important role in all three processes. However, neurons in the ACCs are sensitive to processes that do not conform to the predictions of any of these theories. In particular, these theories cannot account for the processing of rewards both at the time of cues that instruct actions and also at the time that an outcome is delivered. Thus, a better account of the functional properties of the ACC will need to encompass the processing of information during both reward-guided action selection and outcome monitoring.

1.6 A Theoretical framework of ACC function.

Having critiqued three of the most prominent accounts of ACC function, it is apparent that none of them provides a sufficient account of the functional properties of this region. Can anything be learned from the evidence that supports the response selection, error detection and conflict monitoring accounts? Intriguingly, the evidence which tested each of these theories can be interpreted as indicative of two processes being performed within the ACC. Firstly, the ACC is engaged when evaluating information at the time actions are planned and performed, and secondly, the ACC plays an important role in monitoring outcomes and evaluating whether they have achieved their goal. These two processes conform to the basic principles of one of most well founded and prominent theories of learning and decision-making: Reinforcement Learning Theory (RLT) (Sutton and Barto, 1981; Sutton and Barto, 1998). RLT posits that actions are guided by their predicted value at the time of cues that instruct their performance. These values are updated when the outcome of an action reveals that the predicted value was erroneous. Thus, there are clear parallels between the studies that have reported that the ACC is engaged when selecting actions and monitoring their outcomes and the points in time at which RLT framework would predict a response. In this section, an outline is provided of RLT, evidence is presented of how activity in several areas of the brain conforms to the predictions of RLT and then this theory is critiqued as an explanation of the functional properties of the ACCs. In the rest of this chapter, I will highlight a framework

for how the ACCs guides decision-making. Each of the studies in this thesis will then examine whether the same framework can be applied to the ACCg when subjects' monitor the decisions of others.

1.6.1 Reinforcement learning theory

Much behaviour is aimed at maximising reinforcement. To maximise the amount of obtained rewards, one will often have to choose between alternative courses of action, that have differing levels of reward associated with them. Thus, actions (or a series of actions) are assigned an abstract value that is determined by the expectation that they will lead to a beneficial outcome. At the time of making a decision, the value of each course of action is used to guide choices. The value of an action therefore constitutes a prediction about the outcome of the decision to select it. Optimal decision-making is therefore dependent on accurate predictions being made about the value of each course of action. Accurate estimates of value can be learned through the history of reinforcement associated with previous performances of the action. Reinforcement Learning Theory (RLT) is a computational framework that can be used to explain how individuals acquire estimates of the value of actions through experience dependent learning (Schultz, 1998; Sutton and Barto, 1998; Dayan and Balleine, 2002).

RLT has been developed in parallel predominantly by behavioural psychologists and computer scientists, and not neuroscientists (Rescorla and Wagner, 1972; Kawato and Samejima, 2007). However, increasingly it is being used to make predictions and test whether the brain processes information in a manner that conforms to its principles. RLT is underpinned by two important principles (i) a decision is guided by predictions about the value of actions and (ii) when new information reveals that a prediction was erroneous, the values are updated by a prediction error signal (Rescorla and Wagner, 1972; Sutton and Barto, 1998). To illustrate, imagine a monkey learning how much juice will be received for pressing a lever. If the monkey has previously received a small amount of juice following a lever press, it will assign a low value to the action of pressing the lever, predicting that future presses of the lever will lead to only a small juice reward. However, if a large juice reward is received following the next lever press, there is a discrepancy between the monkey's prediction and the actual outcome of the lever press. Thus, the monkey experiences a prediction error. This positive prediction error updates the value of a lever press, increasing its value, making its performance more likely in the future. In contrast, if the monkey performed the lever press and no juice reward was forthcoming, a negative prediction error would occur. As a result the value assigned to the

lever press would be decreased and future performance of the lever press would become less likely.

Although many sophisticated algorithms have been used (Sutton and Barto, 1998), this theoretical framework can also be described using a very simple computational model that was first outlined by Rescorla and Wagner (R-W) (1972). The R-W model assumes that the associative value of an action (or stimulus) changes once new information reveals that the actual outcome of a decision is different from the predicted outcome (Rescorla and Wagner, 1972). Thus, for each performance of an action, its value is updated by a prediction error signal when the outcome reveals that the prediction was erroneous. The evolution of the values for an action are given by:

(1)

Where:

(2)

In both (1) and (2), *n* is the number of times an action, *a* has been performed and η is the rate at which the values are updated. In (1) the value of the action in the future ( ) is a function of current predicted value of the action ( ) added to the prediction error ( ), which is multiplied by the learning rate. The learning rate defines the extent to which the prediction error updates the predicted value. As such, a slow learning rate will minimise the effect of the prediction error and the amount that the value is updated. The prediction error shown in (2), compares the actual value of an action ( ) to the prediction of its value ( ). This discrepancy is what drives the updating of the predicated value in the future. The asymptotic value (λ) of an action is defined as the finite value which can be reached once all learning has occurred and there is no longer a discrepancy between the prediction and the actual outcome. Thus, slow learners will experience a greater number of prediction errors before reaching the asymptotic value than fast learners (see Fig.1.7).

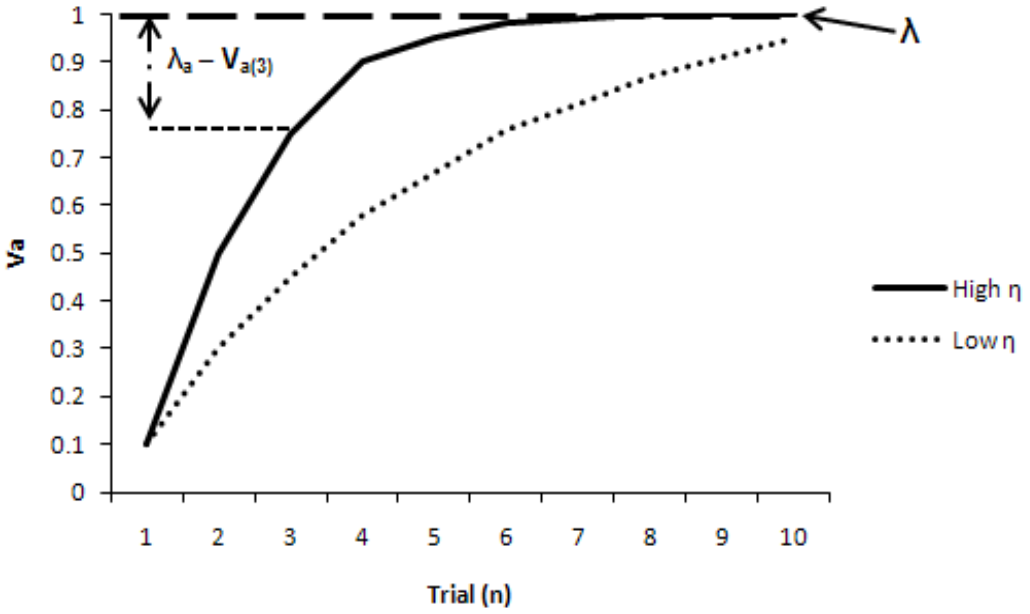

Fig.1.7 Graph illustrating how values are assumed to be learned over repeated performances of an action over trials in the R-W model. Subjects with high learning rates (η) acquire the asymptotic value (λ) in fewer trials than the subject with low learning rates. Over trials the magnitude of the prediction error decreases as the prediction (Va) approaches the asymptotic value.

An important aspect of the R-W algorithm and other models based around the principles of RLT, are the assumptions that it makes about learning. The algorithm includes a free parameter, the learning rate. This free parameter defines the extent to which an individual updates the values of a performed action, regardless of the magnitude of the prediction error. Thus, RLT assumes learning is idiosyncratic, allowing for individual differences in learning and decision-making to be accounted for. In addition, this framework also makes the assumption that learning the value of actions is underpinned by the same computational mechanisms, across individuals and across species.

Another intriguing aspect of RLT is that it proposes that there are areas of the brain that will signal for predictions about the value of actions before they are selected and similarly there will be areas that will signal when these predictions are erroneous. Such an account appears to closely overlap with the common properties of the three theories of ACC function that I referred to in the critique provided above. I suggested that any framework of ACC function must account for activity both when selecting actions and when evaluating their outcomes. Thus, RLT may be a useful a framework in which to discuss the functional properties of the

ACC. However, before I move on to discuss this possibility I will examine the evidence that signals predicted by RLT are processed in any area of the brain.

1.6.2 Reinforcement learning in the brain

RLT was originally developed to explain learning behaviour. In the case of the R-W algorithm, its original purpose was to explain Pavlovian conditioning behaviour of animals (Rescorla and Wagner, 1972). However, this theoretical approach has been extended to explain many different types of learning behaviours, including those that require instrumental learning processes. Whilst such algorithms were designed to examine behaviour, they also make specific predictions about how information is processed and how this relates to learning. As such, these algorithms make predictions about the nature of signals that would be expected in the brain, in order for learning to take place (O'Doherty et al., 2007). As such, specific hypotheses can be tested to examine whether these algorithms not only explain behaviour, but have antecedents in the brain that drive learning.

An important aspect of neurophysiological and neuroimaging methods which investigate reinforcement learning processes, is that they make specific predictions about the parametric nature of responses in the brain (O'Doherty et al., 2007; Dayan and Daw, 2008). In most algorithms that are based on the principles of RLT, predicted values are not binary. As can be seen in the plot above (fig.1.7), the predicted value increases parametrically over repeated performances of an action. Similarly, prediction error signals vary parametrically over learning, decreasing in magnitude as the predicted value approaches the asymptotic value. As such, studies which use single-cell recording methods to investigate the firing properties of individual neurons, examine whether the spike rate of neurons varies with the predicted values or the magnitude of prediction errors. Neuroimaging methods can also be used to examine whether the amplitude of the BOLD response varies with magnitude of predictions or prediction errors.

This thesis is focussed on examining information processing in the ACC and not the whole brain. As such a full discussion of the extensive body of research that has examined the neural antecedents of RLT across the whole brain will not be provided (for comprehensive reviews see Schultz (2006) and Rushworth et al.,(2008)). However, it is important to outline the evidence that highlights the utility of RLT for explaining activity in the brain. Here I will discuss

the most important and relevant findings of neurophysiological and neuroimaging research that show the ubiquity of reinforcement learning signals across different tasks and across different neural systems (O'Doherty et al., 2007; Rushworth et al., 2009).

The earliest work investigating the most significant feature of RLT, the prediction error signal, implicated dopamine neurons in the midbrain as signalling when an unexpected outcome occurs. The firing of neurons in the Ventral Tegmental Area (VTA) varies quantitatively with the magnitude of the prediction error (Schultz, 1998). That is, the greater the discrepancy between a prediction and an outcome, the greater the change in the firing rate of dopamine neurons in the VTA (Fiorillo et al., 2003). These neurons have also been implicated in responding in a directional manner, such that there is an increase in spike frequency for positive prediction errors, a decrease in spike frequency for negative prediction errors and no change when the outcome is the same as the prediction (Schultz et al., 1997; Schultz and Dickinson, 2000). In humans, neuroimaging methods have found reward prediction error signals in other areas including the striatum and the OFC (McClure et al., 2003; O'Doherty et al., 2003; Ramnani et al., 2004b), regions that are connected to the VTA (Williams and Goldman-Rakic, 1998). There is however some debate as to whether the responses in these regions are signed as they are in the VTA, or unsigned, increasing their response with the magnitude of the error regardless of whether it is a positive or negative outcome (Schultz, 2006). In addition, it has also been argued that activity in the OFC may vary with the expected value of an outcome at the time that feedback is delivered and not with the magnitude of a prediction error (Schoenbaum et al., 2009). However, such expected value signals may still be important for prediction error computations that are processed distally (Schoenbaum et al., 2009). Areas connected to the VTA may therefore play an important role in processing information at the time of the outcomes of decisions, particularly when expectations about the outcomes of decisions are violated.

Recent findings however, suggest that prediction error signals may occur outside of systems that process rewards and reflect a rather more ubiquitous functional property of the brain. Predictions error signals are found during a broad range of different processes in neural systems that process very different types of information (Rushworth et al., 2009). Recently, 'predictive coding' models of sensory systems suggest that the discrepancy between a predicted and actual sensory outcome are processed at several levels of sensory systems, continuously updating different types of prediction through error signals (Bar, 2007, 2009). Support for this view has been provided from an fMRI study that has reported activity in several areas within the visual cortex that process discrepancies between predicted and actual

sensory input during a simple forced choice decision-making task (Summerfield et al., 2006; Summerfield and Koechlin, 2008). Den Ouden et al. (den Ouden et al., 2009) examined activity when subjects learnt the association between an auditory tone and the presence of a visual stimulus. They found evidence of prediction error signals in primary visual cortex and the putamen that occurred whenever a visual stimulus was not predicted by the tone. In addition, prediction errors have now been found in other cortical areas such as the pSTS and the paracingulate cortex during social interactions (Behrens et al., 2008; Kennerley and Wallis, 2009a; Burke et al., 2010). Thus, there is tentative evidence to suggest that prediction error signals are found outside of systems that process first-person rewards and reward prediction errors. Indeed prediction error signals may well be an important common mechanism for learning and decision-making in the brain, regardless of the nature of the information that is being learnt.

The other important component of RLT are predicted value signals. Areas which process rewards at the time that actions are selected are found across a large number of cortical and subcortical areas. Indeed, the spike rate of neurons in the Posterior Cingulate Cortex (PCC), ACCs, OFC , intraparietal cortex, the ventral striatum, the VTA and lateral portions of the prefrontal cortex, are sensitive to the magnitude of a reward at the time that prediction are made (Schultz et al., 1997; Schultz, 1998; Schultz et al., 2000; Waelti et al., 2001; Fiorillo et al., 2003; Tobler et al., 2005; Padoa-Schioppa and Assad, 2006, 2008; Kennerley et al., 2009; Kennerley and Wallis, 2009b; Hayden and Platt, 2010; Hillman and Bilkey, 2010; Louie and Glimcher, 2010; Heilbronner et al., 2011; Litt et al., 2011; Pearson et al., 2011). Neuroimaging evidence can be found for each of these areas that supports the notion that these areas are engaged by rewarding stimuli (Ramnani and Miall, 2003; McClure et al., 2004; Rogers et al., 2004; Kable and Glimcher, 2007; D'Ardenne et al., 2008; Rolls et al., 2008; Boorman et al., 2011). Of course it should be noted that each of these areas may be processing value in a different manner or context, as often expected reward values are confounded with other factors, such as reward probability, risk, or motivational salience. However, the claim still stands that these areas are engaged in processing predictions about upcoming events. Thus, much like for prediction error signals, predicted value signals are processed in many different task contexts and in many different cortical and subcortical structures.

There is therefore considerable evidence that there are signals in the brain that conform to the principles of RLT. Moreover, the evidence suggests that RLT may be an important and useful framework for understanding neural and behavioural plasticity in many different contexts. Thus, it is now pertinent to discuss whether there is evidence of predicted values and

prediction errors being processed in the ACC, to examine whether information processing in this area conforms to the principles of RLT.

1.6.3Reinforcement learning theory and the ACC

The connectivity profile of the ACC (particularly the MCC) implicates the area in processing both rewards and prediction error signals. As outlined in section 1.2, this region shows strong projections to the ventral striatum and the OFC, areas that code for both predicted reward values and prediction errors (Kunishio and Haber, 1994; Carmichael and Price, 1995). In addition, it has monosynaptic connections with dopamine neurons in the VTA in the midbrain (Williams and Goldman-Rakic, 1998), which are believed to be the origin of reward prediction error signals in both humans and nonhuman primates (Haber and Fudge, 1997; D'Ardenne et al., 2008). As such, the ACCs receives direct projections from areas which process predictions about rewards and discrepancies between these predictions and actual outcomes.

Over the last decade, the ACC has become increasingly well known for its role in processing rewards. Perhaps the most convincing case for the processing of reward predictions in the ACC comes from single-unit recording studies. Some of the earliest work to identify reward sensitive neurons in the ACC was provided by Shima and Tanji (1998). They trained three monkeys to perform one of two movements (either a push or turn of a handle), whilst they recorded from the rostral portion of the CMA. During separate blocks, one of the movements was rewarded consistently, resulting in the monkeys continuously performing this action trial upon trial. After an extended period, the reward level was lowered, resulting in the monkeys choosing the alternative movement. In addition, monkeys performed control trials where a switch to the alternative movement was cued by an auditory stimulus. They reported four important findings: (i) the CMAr contains neurons for which the spike frequency varies with the level of reward (ii) the CMAr contains neurons that are sensitive to reward information only on the trials immediately preceding the trials where a switch in behaviour was evident (iii) the majority of the recorded neurons responded specifically to the rewarding value of one action or the other (iv) the region of the CMAr that was recorded from was the region which shows strong projections to the primary motor cortex and the spinal cord (V) injection of muscimol, to chemically inactivate the region, resulted in an inability of monkeys to select the more rewarding actions. This intriguing set of results indicated that in the ACCs there are neurons which are sensitive to magnitude of a reinforcer that is associated with a specific

action during decision-making. Since the findings of Shima and Tanji (1998), a large corpus of studies have shown that neurons in area 24c', and particularly the CMAr in its posterior portions, are sensitive to the magnitude of rewards (Shidara and Richmond, 2002; Amiez et al., 2005; Sallet et al., 2007; Quilodran et al., 2008; Kennerley et al., 2009; Kennerley and Wallis, 2009a; Hayden and Platt, 2010; Hillman and Bilkey, 2010; Hayden et al., 2011a). At this point it should be noted that there is also evidence that there are neurons that process reward-related values that are modulated by other factors. However, for the purposes of discussing whether the ACC is engaged by rewards and predictions, I will leave discussion of this research until the next section. As such, neurophysiological evidence supports the view that the ACCs is engaged in processing predictions about the rewarding value of actions.

There is also evidence to suggest that the ACCs is involved in processing predictions that guide decision-making. The performance of monkeys on tasks in which the processing of rewards is a prerequisite for accurate performance, is reduced when the ACCs is lesioned. For instance, Kennerley et al., (2006) showed that lesions to the ACCs of the macaque monkey prevented the ability to sustain continuously reinforced behaviours. Similarly, lesions to the ACCs also impair the ability to learn an association between an action and a reward, but not the ability to learn associations between a stimulus and a reward (Rudebeck et al., 2008). Thus, evidence from animal lesion studies suggests that the ACCs is sensitive to rewarding values and, in particular, learning about the value of actions.

Functional imaging studies in humans support the notion that the ACC is involved in processing predictions about the rewarding value of actions and also that value signals in this area guide choice behaviour. A number of studies report that the ACC is engaged when learning associations between rewards and actions and also when choosing between differently valued options. A recent fMRI study of reversal-learning, where subjects are required to learn that there are switches in the contingencies between one of two actions and one of two rewards, highlighted the importance of the ACC for reward-related learning and decision-making processes. By using a multivariate pattern analysis approach, they were able to decode the choices made by subjects, solely based on the timecourse of activity in three regions: the ACC, the ventral striatum and the supplementary motor areas (Hampton and O'Doherty, 2007). In another study, which specifically examined the processing of predictions and prediction error in the ACC, increased activity was found in the ACCs (in the RCZ) when high reward actions were selected compared to low reward actions in a similar reversal-learning paradigm (Jocham et al., 2009). Studies using different types of reward-based decision-making paradigms, also highlight the ACC as active at the time of cues which are informative as to the value of actions

(Kable and Glimcher, 2007) and also cues that instruct decisions to be made (Rogers et al., 2004), although the activations reported in these studies lay in the anterior portions of the MCC. Thus, there is considerable correspondence between the findings of neurophysiological studies and lesion studies in animals and the findings of neuroimaging studies in humans. Across the sulcus in the MCC there is sensitivity to rewards at the point in time when an individual is able to make predictions about the value of different courses of action. Interestingly, the ACC is also implicated in processing such valuations in conditions where learning and decision-making is required. As such, the functional properties conform to one of the principles of RLT, that predictions are made about the value of actions.

The second, component of RLT is the prediction error. I have already outlined a considerable body of evidence in section 1.6.1 that highlights the ACC as a region that signals when responses are erroneous. However, an important question is whether the ACC signals prediction errors parametrically that update future valuations of actions. There is a considerable body of single-unit recording studies that have tested this hypothesis. Perhaps the most convincing investigation into reward prediction error signals in the ACC, was performed by Matsumoto et al., (2007). In their task, monkeys learnt arbitrary associations between actions and secondary reinforcers (visual cued positive or negative feedback cues) by trial and error. They found neurons in the ACCs (in the MCC, lying rostral to the CMAr) that responded to negative feedback when an incorrect action was performed and a separate set of neurons that responded to positive feedback when a correct action was performed. Crucially, these neurons showed their greatest response on the first presentation of a type of feedback for a particular action, with decreased responses for each presentation of the same feedback for that action i.e. neurons that responded to negative feedback showed the greatest response at the time of the outcomes of the first performance of an erroneous action, with a decreasing response each time that action was subsequently performed. Similarly, positive feedback preferring neurons showed the greatest response to the first correct feedback for an action and then decreasing responses at the time of the feedback on subsequent trials where the correct action was performed again. Thus, these neurons showed a profile of response that closely matched what would be predicted by a reinforcement learning algorithm, with decreasing prediction error responses occurring as the discrepancy between a prediction and outcome becomes smaller. Other studies also support this claim. Amiez et al., (2005) identified neurons in the ACCs, including neurons within the CMAr and some slightly more rostral in the sulcus, that signalled when erroneous responses were made. However, the response in these neurons was modulated by the magnitude of reward that was predicted. That is, the greater

the predicted reward, the greater the response in ACCs neurons when an action was erroneous. Such results also support the assertion that there are neurons in the ACCs that signal the degree of discrepancy between a prediction and an outcome. However, crucial for the RLT framework, is that these prediction error signals result in a change in the value that is assigned to an action. One study directly tested this claim. Rhesus monkeys performed a zero-sum game, making binary choices. As in the other studies, they found neurons, in the CMAr and a portion of the sulcus that lay rostral to it, that responded to unexpected outcomes and also neurons that responded to predictions about the value of actions. However, they found that activity in these neurons was modulated by the history of previous rewards associated with a choice, in a manner that conformed to the predictions of a reinforcement learning algorithm (Seo and Lee, 2007).

The neurophysiological evidence therefore indicates that there are neurons extending across the CMAr and more rostral portions of the MCC, which respond in a manner that would be predicted by RLT. However, is the firing of single neurons, informative for the function of the ACC as a whole? Neuroimaging and electrophysiological studies suggest that prediction errors are an important functional property of the ACCs. In section 1.5, I reported the results of a host of studies that have implicated the ACC as the source of the ERN, a signal which occurs when erroneous responses are made. Previously many proponents of the error detection theory claimed that the ERN only occurs when feedback reveals that an erroneous response has been performed. However, some studies have suggested that actually these signals closely match those that are found using single-celled recording in monkeys. To examine this issue Holroyd et al., (2009) performed three studies in which they examined the magnitude of the ERN. They showed that the ERN occurs at time of feedback that signals either a positive or negative prediction error, and shows the greatest amplitude when future behaviour is dependent on that feedback. Thus, these results support the view that the ACCs codes prediction error signals at the time of the outcomes of decisions.

Whilst the cases provided from neurophysiology and electrophysiological accounts highlight an important role for the ACCs in processing prediction errors, the evidence from fMRI studies is less supportive. Studies which have examined activity in the brains of subjects who are performing tasks requiring reinforcement learning mechanisms, have not found activity in the ACCs which varies with the prediction error parameter from a RLT algorithm (Behrens et al., 2007; Brovelli et al., 2008; Jocham et al., 2009). For example, Brovelli et al. (2008) fitted a R-W algorithm to the behaviour of subjects in a conditional motor learning paradigm (i.e. learning arbitrary associations between instruction cues, actions and secondary reinforcers). Whilst

they identified activity in the ventral striatum that varied with the prediction error parameter, no such activity was found in the ACC. In contrast to such findings, the number of studies that report ACC activity when the outcome of a decision is unexpected is considerable. This includes studies that show that the ACCs signals both unexpectedly rewarded and unexpectedly unrewarded actions (Jessup et al., 2010; Nee et al., 2011). More recently, some have suggested that the ACC may be involved in the reinforcement learning process, but may code the learning rate rather than the prediction error. When learning the value of actions, where the probability of receiving a reward is uncertain, decision-making can only be optimised if individuals adapt their learning rate to reflect the volatility of the outcomes of decisions. Behrens and colleagues (2007) found that activity in the ACCs (putatively in the MCC) varied with the learning rate in a reinforcement learning model, when the volatility of associations between decisions and outcomes was variable. However, a very recent paper has countered this claim, suggesting that the ACCs may signal prediction errors when several actions (and therefore predictions) are embedded within a larger task goal. Ribas-Fernandes et al. (2011) used a hierarchical reinforcement learning model, in which the predictions and prediction errors for several actions (or subroutines) were embedded in a larger task goal. They found that activity in the ACCs (putatively in the MCC), in a region extending over the sulcus, varied with the prediction error for each of the embedded actions. Thus, the fMRI literature paints a slightly confusing picture, although there is clear evidence that in certain tasks contexts, the BOLD signal varies with the discrepancy between a predicted and actual outcome of an action.

In summary, there is considerable neurophysiological, electrophysiological and neuroimaging evidence that supports the view that the ACC is engaged in processing predictions about the value of actions. In addition, although the fMRI evidence only provides tentative support, there is also evidence that the ACCs signals when there is a discrepancy between the predicted value of an action and its actual value, at the time of an outcome of a decision. I therefore argue that such evidence supports the notion that the ACCs codes for the value of actions and updates these values in a manner that conforms to the principles of RLT. In this thesis, I will test whether the ACCg conforms to the same computational principles when monitoring the actions and decisions of others.

1.6.4 Rewards, values and discounting variables

The value of a reward is dictated by more than just the magnitude of a reinforcer. Reward values are discounted by other factors (or variables), making the performance of the actions associated with the reward less likely. As such, actions can be assigned a net value that reflects the amount of reward discounted by its associated costs. The probability that an action will result in a reward, the effort required for a rewards receipt, and the length of time before a reinforcer is delivered, are all important factors that modulate reward values and impact upon decision-making. For example, a monkey may forego the receipt of a large volume of fruit at a distal location, in order to receive the smaller amount of reward at a proximal location (Hayden et al., 2011a). The effort (in this example a metabolic cost) of travelling the greater distance has discounted the value of the larger reward, such that the reward that is smaller in magnitude is preferred (Phillips et al., 2007). Examples of how effort, probability and temporal delays can impact upon decision-making can be found in many different species. Understanding the neural mechanisms that underpin the process of discounting the value of rewards is, therefore unsurprisingly, a topic that has attracted a considerable amount of research interest. Several areas have been implicated in processing the variables that discount the value of rewards, including the ACCs.

An important body of work by Kennerley and colleagues, reported in two papers (Kennerley et al., 2009; Kennerley and Wallis, 2009a), shows strong evidence in support of the claim that the ACCs processes variables that discount reward values and guide decision-making. Kennerley and his colleagues taught monkeys the associations between a series of abstract cues, actions and rewards. Each cue was associated with either (i) a specific number of lever presses (manipulating effort), (ii) a specific amount of juice (manipulating reward magnitude) and (iii) a specific probability of obtaining a juice reward. Five different levels of effort, reward magnitude and reward probability were manipulated and each cue was associated with one of each variable. By manipulating these variables they were also able to examine whether neurons multiplex information about each of the variables. Subjects were required to make choices between a pair of these stimuli on each trial. They recorded from neurons in the ACCs (which putatively extended across the CMAr and anterior parts of the MCC) both at the time of instruction cues and also at the time the outcome of the decision was presented. They reported several important findings: (i) there are neurons in which the spike frequency varies with the level of the different variables that they manipulated (effort, probability and reward magnitude), (ii) there are neurons that multiplex the different variables signalling the net value

of rewards, i.e., less effort for, or higher probability of, a reward equates to greater net value and increased firing in these neurons, (iii) some neurons signal only at the time the instruction cue is presented and some only at the time at the time of outcome, (iv) other neurons code for expected values at the time of the instruction cue and then actual values at the time of the outcome. This research therefore provided convincing evidence that the ACCs processes the variables that discount the value of rewards. Moreover, it suggested that the ACCs process reward magnitudes discounted by the cost of performing the actions. Another recent study showed that neurons in the ACCs are important for processing variables that guide complex foraging decisions (Hayden et al., 2011a). The spike rate of neurons in the ACCs steadily increased the longer spent in a food patch until the monkey made the choice to move to a new food patch. Thus, it appears that neurons in the ACCs were coding for the motivation for moving to a new foraging patch. As the resources in the current location are increasingly depleted, the cost of moving to a new location decreases and the value of moving increases. This therefore suggests that the ACCs plays a key role for assigning values to different courses of actions during complex decision-making.

This neurophysiological research therefore suggests that across the sulcal MCC, there are neurons which multiplex information about rewards and costs, providing the net value for a course of action. Support for this claim can be found in the neuroimaging literature. A number of studies have shown that activity in the Rostral Cingulate Zone (RCZ), believed to be the homologue of CMAr, varies with the probability a reward will be received and also with the magnitude of a reward (Rogers et al., 2004; Rolls et al., 2008; Bickel et al., 2009; Wunderlich et al., 2009). Both of these variables modulate the firing of neurons in the CMAr (Shidara and Richmond, 2002; Sallet et al., 2007; Quilodran et al., 2008). In addition, Croxson et al., (2009) performed a study which examined activity at the time of cues that indicated the magnitude of a reward available and also the number of actions (or effort) that would be required for its receipt. They showed that activity in the ACCs signalled an interaction between the level of effort and reward, with the highest activity occurring when the cues signalled that a high reward was available with only a small amount of effort was required for its receipt. Other neuroimaging studies have also shown that activity in the ACCs varies with value of rewards discounted by the length of delay before their receipt (Kable and Glimcher, 2007; Luhmann et al., 2008; Peters and Buchel, 2009; Prevost et al., 2010). The study by Behrens et al. (2007) that has already been discussed also reported activity in the ACCs that coded for the volatility of the history of rewards associated with an action. In this study, volatility reflected an interaction between the reward probability and stochasticity of the reward schedule over the

history of trials. Thus, there is neuroimaging evidence that supports the claim that the ACC codes variables which guide decision-making by discounting the value of rewards associated with an action.

In summary, there is considerable evidence from lesion, neurophysiology, and neuroimaging studies that highlight the ACCs as an area that as important for processing the variables that guide action selection and decision-making. This area codes for the motivation for the performance of an action in the form of a prediction about the value of performing it. These values can sometimes reflect more than one variable that can guide decision-making, such that the magnitudes of rewards are discounted by associated costs. In addition the evidence also suggests that predictions about the value of an action are updated through prediction error signals in a manner that conforms to the principles of RLT.

1.7 Statement of Thesis Aims

At the beginning of this chapter several claims are made about the contribution of the ACC to social cognition. Firstly, it is argued based on lesion studies and anatomical evidence, that it is the gyral surface of the MCC (which is referred to as the ACCg) that is involved in the processing of social information and not the adjacent sulcus. Secondly, a small number of studies are used as evidence to suggest that the gyrus may process similar information to the sulcus, but process this information about others, rather than about oneself. Due to an absence of a theoretical account of the contribution of the ACCg to social cognition, it is therefore suggested that important clues as to how it may contribute may be found by reviewing the literature investigating the functional properties of the ACCs. A framework is outlined for information processing in the ACCs. This framework proposes that the ACCs conforms to the principles of RLT, processing predictions about the values of one's own actions and updating values through prediction error signals. In addition, in this framework it has been suggested that the ACCs processes the variables that discount the value of rewards, such that activity in the ACCs signals the net value (or motivation) for performing actions. Does the ACCg process others' predictions about the discounted value of their actions and also signal when their predictions are erroneous?

The aim of this thesis is to investigate whether the ACCg processes the same information about the actions of others, that the ACCs processes about one's own actions. Specifically two important questions are asked about information processing in the ACCg:

(i) Does information processing in the ACCg conform to the principles of RLT, signalling the predicted value of others' actions and also signalling when others' predictions are erroneous?

(ii) Does the ACCg process the discounted value of others' actions, in the same manner that the ACCs processes the discounted value of first-person actions?

This thesis contains four behavioural and fMRI experiments examining information processing in the ACC during social decision-making tasks. Each of these experiments uses an event-related design to examine activity time-locked to stimuli that signal the value of actions, or that signal that predictions are erroneous. The studies use a combination of both factorial experimental designs and computational modelling approaches to examine the two core themes outlined above. The four key questions that are addressed by each experiments are:

1. Do prediction error signals in the ACCg code for others' false-beliefs about the outcomes of their decisions?
2. Do the predictions error signals in the ACCg conform to the computational principles of RLT?
3. Does the ACCg process others' predictions about the net value of effortful actions?
4. Does the ACCg process the value of delayed rewards when they are discounted in a manner that conforms to others' preferences?

Chapter 2: General Methods

Within this thesis, four experiments using both behavioural and functional Magnetic Resonance Imaging (fMRI) methods will be reported. In this chapter, some of the methods which are generic to several of the experiments will be outlined.

2.1 Procedures

2.1.1 Subjects

The subjects for all of the experiments were recruited from the Undergraduate and Postgraduate student population at Royal Holloway, University of London. All subjects were healthy, screened for psychological and neurological conditions, right-handed and within the age range of 18-35. All participants gave written informed consent. The studies were approved by the Royal Holloway, University of London Psychology Department Ethics Committee and conformed to the regulations set out in the CUBIC MRI rules of Operations (http://www.pc.rhul.ac.uk/sites/cubic/). The experiments in this thesis involved deception, thus secondary consent was required following data collection, ensuring that subjects still consented to the use of their data. Debriefing sheets are included in appendix A.

2.1.2 Piloting and Quality Assurance

Before fMRI data were collected for experiments, several important steps were performed to minimise noise in the fMRI data and to ensure that subjects performed the task in the manner intended by the experimental design. Behavioural pilot experiments were conducted prior to the full experiments inside a mock fMRI scanner. For these behavioural pilot experiments, it was essential that subjects performed the tasks in the manner intended by the experimental designs. This included both the performance of the task itself and belief in the deception (see 'social decision making task procedures' below). If problems were identified in this pilot, then elements of the experiments were altered and a new behavioural pilot study was conducted. Once a behavioural pilot was successful, quality assurance (QA) was conducted on the Echo-planar image sequence that would be used on the MRI scanner. This involved scanning a 'phantom' (an anthropogenic Perspex sphere containing Magnetic resonance sensitive

material) using the same scanner parameters that would be used during the experiment. Several measures were used to assess this data. Firstly, a slice-by-slice Fourier Transform was used to look for any cyclical changes in signal intensity across the scanning period. Secondly, plots of the signal intensity (corrected by the mean of each slice) over volumes, the mean volume signal intensity in each volume and the standard deviation of each volume were used to determine whether there were any spikes (i.e volumes in which all slices showed a significantly different signal intensity to the mean and to neighbouring volumes). Thirdly, mean, standard deviation and signal to noise images were created to examine the severity of any aretfacts such as ghosting and any other transient artefacts. Experimental data collection only commenced if there were no spikes, no ghosting effects and no transient noise artefacts. The results of the data collected before each experiment are included in the Appendix (B).

2.1.3 Social Decision-Making Task Procedures

Each of the experiments in this thesis required subjects to be deceived in order to maintain experimental control. For three of these studies (reported in chapters three, four and five), the subjects were instructed that they would perform the task whilst inside the MRI scanner, with another participant performing the task in the control room next to the scanner. The scanned subjects were informed that they would be monitoring the responses of the other participant in real-time. The subjects were also informed that this other participant was being paid for their participation. However, in each of these studies, the responses observed by the scanned subject were in fact pre-programmed, computer-controlled responses (details are provided in each chapter about the nature of the responses they observed). In addition, the participants situated outside the scanner were actually confederates. The confederates took on a different pseudonym for each participant. This avoided the possibility that subjects would become aware of the deception by conversing with other subjects after training but before scanning (as they were conducted on separate days). In addition, to remove potential confounds that may surround differences in the interpersonal interactions between subjects, it was ensured that the confederates were not acquainted with the scanned subjects.

In the experiment reported in chapter 6, the nature of the deception differed from the other studies. The subjects were not engaged in a social interaction with a confederate. Instead, during a training session they were told that they were learning what was the majority behaviour on the task. They then reproduced this behaviour once inside the scanner. However,

the behaviour they learnt during the training session was not that of a real majority behaviour but invented for the purposes of the experiment.

In each of these studies, the deception was crucial to the experimental design. In order to elucidate whether subjects had maintained a belief in the deception throughout each experiment, subjects were asked a series of standardised questions immediately following data collection (see appendix C). In each experiment, before the subjects were debriefed, they were asked questions that might provoke them to indicate that they had not maintained a belief in the deception throughout the experiment. Following debriefing, in which they were informed of the nature of the deception, they were asked a series of questions to determine whether they had been deceived throughout the experiment. If a subject reported being aware of the deception during the experiment they were excluded from the behavioural and fMRI analyses. This approach has previously been used in other social decision-making experiments to ensure that subjects perform the task in the manner intended by the experimental design (Ramnani and Miall, 2004; Apps et al., accepted).

2.2 Apparatus

All of the behavioural and fMRI data reported in this thesis was collected using the same experimental setup and apparatus (see fig.2.1 for a schematic). Identical setups were used in the MRI scanner and also in a mock MRI scanner that was used for both training participants before scanning and also for collection of pilot behavioural data. All subjects were scanned using the 3tesla Siemens Trio scanner housed at Royal Holloway, University of London. During data collection, subjects lay supine in the scanner with the fingers of the right hand positioned on a on MRI-compatible 4-button response box. Stimuli were projected onto a screen behind the subject and viewed in a mirror positioned over the subject's head. Presentation software (Neurobehavioral Systems, Inc., USA) was used for experimental control (stimulus presentation and response collection). A custom-built parallel port interface connected to the PC running Presentation received transistor-transistor logic (TTL) pulse inputs from the response keypad. It also received TTL pulses from the MRI scanner at the onset of each volume acquisition, allowing events in the experiment to become precisely synchronized with the onset of each volume acquisition (in the mock MRI scanner, these were simulated by the Spike2 software (see below)). The timings of all events (stimulus presentation, button presses and TTL pulses from the scanner) in the experiment were sampled accurately, continuously and

simultaneously (independently of Presentation) at a frequency of 1 kHz using an A/D 1401 unit (Cambridge Electronic Design, UK). Spike2 software was used to create a temporal record of these events. Behavioural data analysis was performed offline, and event timings were prepared for subsequent general linear model (GLM) analyses of fMRI data. Analyses of fMRI data were conducted in SPM5 (chapter three) and SPM8 (chapters four, five and six) on a Viglen Genie machine running 2x Intel(R) 64MHz duo core 4500 CPUs, with 1.9GB or RAM, running Linux (Ubuntu) and Matlab (2007a; MathWorks Inc.). The computational modelling aspects of the thesis were scripted and analysed in Matlab. The statistical analysis of behavioural data was conducted in SPSS 12 and 14 on a Viglen machine with 2x Intel Core 2 6320 CPU's with 1.97GB RAM running Windows XP (SP3).

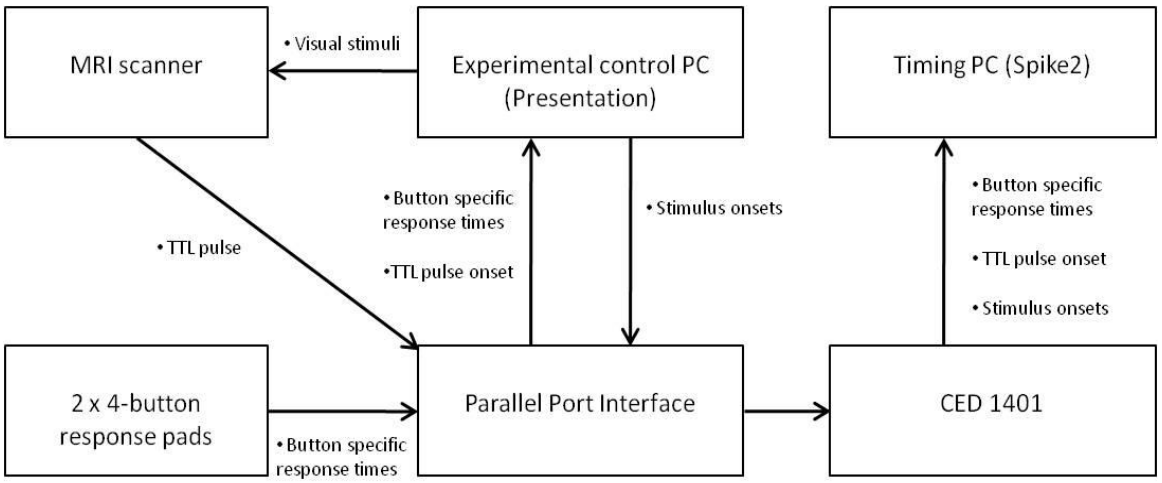

Fig.2.1 Apparatus used for fMRI studies. A similar setup was used inside a mock MRI scanner, for pilot experiments and training sessions. In the mock MRI scanner the TTL pulses were mimicked by Spike 2.

2.3 Data Acquisition

2.3.1 MRI

All of the experiments in this thesis involved the collection of both structural and functional images. In each experiment, the first image acquired was a high resolution T1-weighted structural image, acquired at a resolution of 1x1x1mm using an MPRAGE sequence. These anatomical images were used for the purposes of normalization (see below), the creation of masks used for correcting for multiple comparisons (see 'anatomical localization and small volume correction') and displaying results.

Following the acquisition of structural images, subjects performed the experimental tasks during the acquisition of Gradient Echo-planar Images (EPI). EPI imaging has become the most commonly employed technique in human fMRI, due to the sensitivity to changes in the Blood-Oxygen-Level Dependent (BOLD) signal (Ogawa et al., 1990) and rapid acquisition time (a whole brain volume can be acquired following a single Radio-Frequency pulse) (Mansfield, 1984; Stehling et al., 1991).

Despite their utility, EPI images are susceptible to losses in signal and image distortions. These are particularly prominent at borders between substances which have different magnetic properties, such as the borders between differing tissue types and the borders between tissue and air. At these borders, inhomogeneities in the magnetic field can occur (known as susceptibility artefact), distorting the intensity of the voxels in the image. This is particularly problematic for studies investigating areas of the brain which lie close to air/tissue borders and particularly for examinations of subgenual, orbitofrontal and frontal polar cortices (Lipschutz et al., 2001). A number of methods have been proposed to reduce susceptibility artefact. One possible method is to use tailored, subject-specific radio-frequency pulses for excitation (Stenger et al., 2000). However, such an approach is impractical for multiple subject experiments of the nature employed in this thesis, as it requires the design of a unique excitation pulse for each subject before the start of the fMRI experiment. An alternative approach may be Z-shimming, in which multiple images are acquired with different gradient pulses (Glover, 1999). However, the acquisition of multiple images significantly reduces temporal resolution, making this approach unsuitable for event-related designs such as those reported here. Alternative methods, such as placing resistive shim coils in the mouths of subjects (Hsu and Glover, 2005), were also not employed as they were not practical.

Two methods were employed to minimize the impact of susceptibility artefact on the data collected in this thesis. Firstly, images were acquired at an oblique angle to the AC-PC line. Whilst susceptibility artefact is a result of field inhomogeneities caused by differences between the properties of neighbouring tissues, susceptibility is also influenced by in-plane and through-plane gradient susceptibility i.e. different gradients in the image have an increased susceptibility to signal loss. This loss of signal can be redistributed to areas which are not as susceptible to dropout by acquiring images at a different orientation (Merboldt et al., 2001; Deichmann et al., 2002). In this thesis, hypotheses are being tested relating to the processing of rewards in the ACC. It was therefore essential to examine activity in areas which have strong connections to the ACC and that are also well-known for the processing of rewarding stimuli. One area which shows such a profile is the Orbitofrontal Cortex (OFC) (Ongur and Price, 2000; Rolls, 2000; Schultz, 2002). The OFC is one area where considerable signal loss can occur due to susceptibility artefact. It was therefore important to ensure that the slice orientation used in each of the studies reduced the impact of susceptibility artefact. Previously it has been shown that using an orientation of 30˚ oblique to the AC-PC line decreases signal loss and increases sensitivity to the BOLD signal in the OFC (Deichmann et al., 2003). This hypothesis was tested when using the EPI parameters that would be used during the experiments contained in this thesis. 6 EPI images were taken at each of four angles (0˚, 25˚, 30˚ and 40˚) and two different in-plane acquisition directions (Anterior-Posterior or Posterior-Anterior). The results confirmed the hypothesis that an angle of approximately 30˚ oblique to the AC-PC line showed increased signal in the OFC and also showed increased signal in the temporal lobes, cerebellum and in the striatum. As a result, for all of the studies in this thesis, subjects were scanned with an angle of orientation as close to 30˚as possible. In some cases this angle was slightly less than 30˚ as it would not have been possible to scan the whole cortex at that angle. However, it was found that the lowest OFC signal occurred when images were acquired at 0˚. As such, acquiring at angles greater than 0˚, but less than 30˚ may be beneficial compared to acquiring parallel to the AC-PC line. It should also be noted that the position of each subject's head differs once inside the scanner. The angle of the AC-PC line with respect to the orientation of the field of view therefore differs from subject to subject. The angle of acquisition of the AC-PC line was therefore approximated and for this reason the angle of acquisition was never exactly 30˚.

The second approach used to minimize artefactual effects was to unwarp the images post-hoc using fieldmaps. The rationale and implementation of this is discussed in the 'Realignment and unwarping' section below.

2.4 fMRI Analysis

2.4.1 Pre-processing

2.4.1.1 Realignment

Changes in the signal intensity of one voxel over time can occur as a result of subject head motion. Such changes in signal intensity can be much greater than those resulting from task-related changes in the BOLD response. As a result, head motion significantly increases the size of the error term in a General Linear Model (GLM) and in some cases can explain as much as 90% of the variance in an analysis (Friston et al., 1996). This significantly decreases the sensitivity of the analysis to the conditions of interest. This is even more of a concern when head movements are temporally correlated with experimental conditions. If this occurs for some experimental conditions, but not all, head motion drastically increases the chances of there being arefactual activations, particularly on the cortical surface. Such activations would lead to false rejections of the null hypothesis (a type I error). It is therefore essential to control and correct for head motion as much as possible.

In SPM (both SPM5 and SPM8) realignment takes place in two stages. Firstly, registration, in which the amount of movement is calculated by comparing the first image acquired to all subsequent images. Registration is performed using a rigid body transformation that minimises the differences between a reference scan (a mean image or the first image acquired) and source images (i.e each successive EPI scan). Three translations (x, y and z directions) and three rotations (roll, pitch and yaw) are estimated using a least- squares cost function, which minimises the sum of squared differences between the source images and the reference image. This least squares cost function selects the optimal set of translations and rotations to minimise the differences between the reference and source images, and therefore optimally explain the data (Friston et al., 1995a). Once the images have been registered to a reference image, they are then transformed (or resliced) based on the parameters estimated during registration. The image transformation occurs through interpolation, where the intensity in the transformed image is determined based on the intensity in the original image. Image intensity values at a given voxel are taken from the nearest neighbours as a basis for the interpolation.

Although the process of realignment significantly reduces the effects of head motion (Friston et al., 1996), it is only an approximation and there is still considerable variance that cannot be

explained due to the linear nature of the correction. As such, noise which cannot be explained by linear correction methods can also confound statistical results as a result: (i) movements between slice acquisitions (ii) Interpolation artefacts (Grootoonk et al., 2000) (iii) spin-excitation history effects (when the current signal becomes a function of movement history) (Friston et al., 1996). To reduce some of these effects, the head motion parameters estimated during registration were included in the GLM as regressors of no interest. Although, due to the non-linear nature of these artefacts, this approach is unlikely to eliminate them completely.

2.4.1.2 Realignment and Unwarping

As stated above, susceptibility artefact and head motion can independently introduce noise into fMRI data and analyses. However, an additional problem is the interaction between these two sources of noise (Andersson et al., 2001). When a subject moves, different points in the EPI image become distorted due to susceptibility aretefact. The result is an apparent change in shape of the subject's brain that depends on head motion. Thus, the linear method of registering images described above becomes less accurate, as the images have different geometric properties. However, it is possible to use field maps, which provide a description of the properties of the magnetic field, to undistort and correct for inhomogeneities in the magnetic field. By combining the field maps with the EPI images, a voxel displacement map can be created which characterises the distortions that are common to all images in each subject. It is also possible to calculate the distortions between images and use the resulting parameters to undistort the EPI images after realignment. By undistorting images in this manner, it is possible to improve realignment (Andersson et al., 2001) and as a result improve normalisation (Hutton et al., 2002; Cusack et al., 2003) and group level statistical power (Cusack et al., 2003). This approach was applied in chapters four, five and six.

2.4.1.3 Normalization and Unified Segmentation

The aim of the fMRI studies in this thesis is to make inferences about information processing in the population. In order to make such inferences, one tests hypotheses about the nature of information processing in one anatomical location in all subjects, which are a sample of the population. One therefore has to ensure that one voxel reflects the activity of the same

anatomical location in all subjects. To achieve this, EPI images are registered with a template which is in a standard space. Thus, the brains of all subjects in fMRI experiments are approximately within the same stereotactic space and coordinate system. In this thesis images are warped into the same space as Montreal Neurological Institute (MNI) template which is recognised by the International consortium for Brain Mapping as the standard template.

The conventional approach to normalisation is to warp the images directly into the space of the template image (the approach applied in chapter three). However, in this thesis, an improved method is used in which the warping is constrained and optimised by information from the subjects anatomical images. The anatomical images are first segmented into separate grey and white matter images based on template grey and white matter images.

In this thesis, anatomical images were segmented using a unified segmentation procedure (Ashburner and Friston, 2005) in which structural scans are segmented, spatially normalised and bias corrected within a single generative model. The anatomical image is passed through an affine registration, in order to align the images to the template. Tissue probability maps, which are based on the probability of a particular voxel being grey or white matter in the population, are passed through a deformation algorithm, in order to improve the fit to the registered anatomical image. This process is optimised by minimising the sum of two terms. The first gives the probability of the data given the warping parameters and the second gives the probability of the parameters, which becomes lower when the deformations are unlikely. This is implemented using Bayesian principles. In addition, the images are corrected to remove any smooth, spatially varying, MR physics related bias. The outputs of this process are bias corrected, subject-specific segmented grey and white matter probability maps, which are in the ICBM template space. The parameters from this generative model can then be used to warp all the subjects' EPI images into the template space. This approach has been shown to result in greater overlap of anatomical structures between one brain and the template, as well as across subjects, with a particular improvement shown in the ACCg (Klein et al., 2009). This greater overlap affords greater statistical power to the GLM analysis, as each voxel is more likely to represent the activity of the same anatomical location across subjects.

2.4.1.4 Spatial Smoothing

When performing a multi-subject fMRI experiment, one makes important assumptions that must be met for statistical inferences to be valid. In particular, it is assumed that errors are spatially normally distributed and that the error terms are a representation of a smooth, underlying Gaussian field (Friston et al., 1995b). In order to meet these assumptions, a 3D Gaussian kernel is applied to each voxel, smoothing the intensity values so that they are normally distributed. The optimal kernel corresponds to the size of the effect that is expected, which is assumed to be roughly 5mm. However, when averaging across subjects, it is often necessary to smooth using a larger kernel (8mm is the kernel applied in this thesis), in order that homologies between different subjects' functional anatomy are expressed. It is important to note that although smoothing increases the validity of the statistics and also increases the signal to noise ratio, it is at the cost of some of the spatial resolution which is afforded by fMRI. Once smoothing has been completed the images can be statistically analysed

In summary, the EPI images acquired in chapters four, five and six of this thesis were preprocessed, using unwarping, realignment, normalisation and smoothing. The pipeline used in these studies is outlined in figure 2.2.

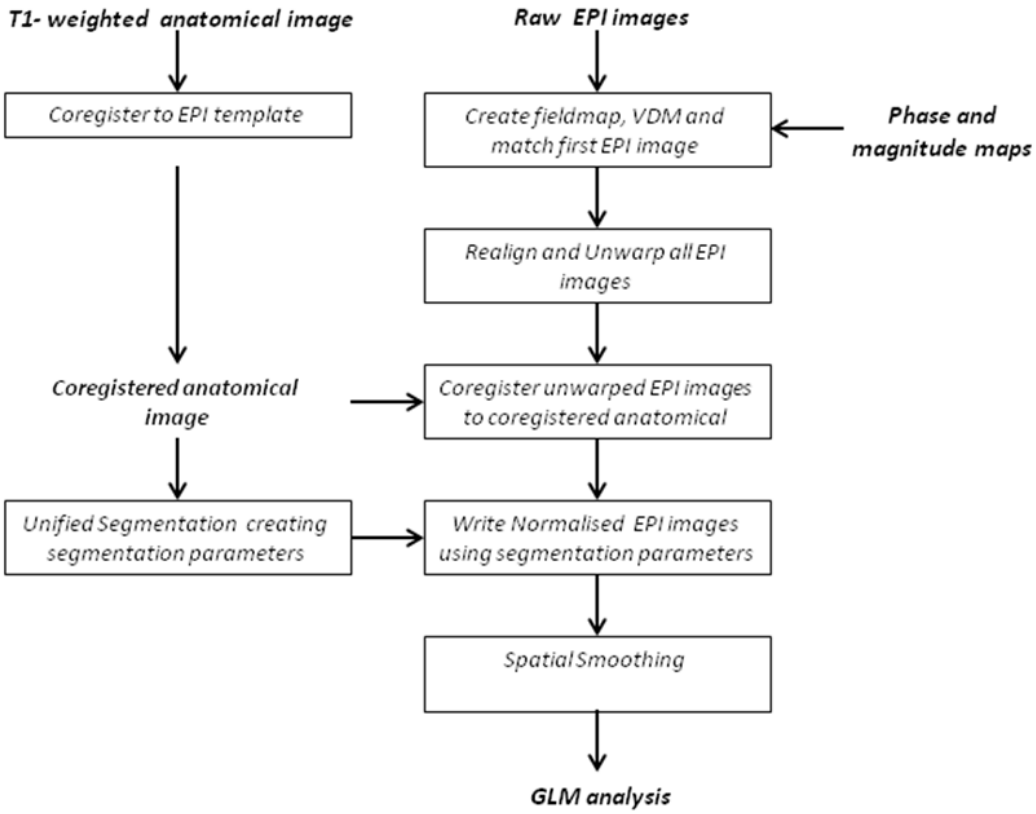

Fig.2.2 Outline of the pre-processing pipeline implemented in chapters four, five and six.

2.4.2 Statistical Inference

2.4.2.1 The General Linear Model

Statistical Parametric Mapping is typically the term given to the use of a General Linear Model (GLM) and Gaussian Random Field Theory (GRF) to produce spatially extended maps of univariate statistical tests conducted on fMRI data. This approach involves the application of a mass-univariate GLM, at each and every voxel in preprocessed EPI images. The GLM is used to partition the observed time-series in a given voxel (Y) into variables of interest and confounds (X), scaled by an estimated parameter (β), plus a noise or error term (ε). Thus, mathematically the GLM represents the data at a given voxel as:

Or:

$$= X\beta + \varepsilon$$

The GLM is implemented in SPM through the creation of a design matrix (X). The design matrix represents the time-course of the experiment as experienced by the subject. Each variable in X represents one experimental condition or one confound that explains some event (or events) that took place during the experiment. These explanatory variables are regressed against the data, i.e. the β parameter is estimated to examine the extent to which the data Y can be explained by each of the explanatory variables.

In fMRI the BOLD response (the dependent variable) poses a problem for a typical GLM analysis, due to its delayed timescourse. The BOLD response is temporally delayed with respect to the neural response following stimulus presentation, typically peaking at ~6s, undershooting after ~12s and returning to baseline only after ~32s. If events were modelled as simple stick functions at the time of stimulus onset in a GLM, the analysis would not be able to model the BOLD signal following a stimulus. Thus, the onsets of stimulus presentations are convolved with a basis function, which is a prediction of the response that would be expected following a stimulus event. Thus, in a given voxel, the predictions of an experimental variable in the design matrix X, should closely match the BOLD signal Y, if there is stimulus induced activity. The most typically used basis function and the default in SPM is the canonical haemodynamic response function (HRF). The HRF is the sum of two gamma functions which,

following a stimulus event, model a peak 6s after onset and an undershooting at 12s. In this thesis, the canonical HRF was used in several chapters (4, 5 and 6).

An important aspect of this thesis was the use of computational models to examine both behaviour and the BOLD response. When using such computational approaches it is assumed that the value of a parameter in the behavioural model will correspond to a parametric change in the magnitude of the BOLD response. In SPM, one can therefore use parameters from a computational model as modulators of an experimental variable. When using a parameter as a first-order modulation, the convolved HRF is scaled such that the height of the peak corresponds to the magnitude of the scaling parameter. This creates a new regressor, which models the same points in time during the experiment as the original explanatory variable. However, the new regressor makes different predictions about the magnitude of the BOLD response following each stimulus onset.

2.4.2.2 Temporal Filtering

Temporal filtering occurs in two forms in a GLM analysis. Firstly, to remove low-frequency noise. fMRI time-series are dominated by low-frequency signals which act as noise and reduce the efficiency of statistical tests. This noise comes from several sources, including changes in ambient temperature ('scanner drift') and physiological sources such as breathing and heart rate. In order to reduce the impact of this noise, a high-pass filter is used (typically a 128s/0.008Hz filter). This removes some of the effects of the low-frequency noise on the data.

Temporal filtering is also used to ensure that data does not violate assumptions about sphericity. As stated above, the use of GLM relies on the assumption that error is Gaussian. However, the time-course of the BOLD response is slow and EPI images are acquired sequentially, resulting in temporal autocorrelations in the images, which violate assumptions necessary for parametric statistical tests. It is important to correct for these autocorrelations by prewhitening the design matrix (Woolrich et al., 2001; Smith et al., 2007). That is making sure that each sample of the BOLD response in a given voxel is independent or "white". In SPM, the design matrix is prewhitened using autoregressive error (Z_t) and white noise(e_t) terms, which model the temporal correlations and the nonspherciity of the data within the GLM as follows:

$$Y_t = X_t \beta + Z_t + e_t$$

In SPM, following prewhitening, the parameters (beta coefficients) in the design matrix are estimated, allowing for statistical tests to be conducted and inferences to be made.

2.4.2.3 Statistical Parametric Mapping and Inference

In fMRI the most common statistical approach is to make post-hoc contrasts at each and every voxel, between regressors in the GLM. These contrasts take the form of either a student's *t*-test (or *t*-contrast) or an ANOVA (*F*-contrast). The result of these contrasts is a spatial map of parametric statistical tests (either SPM{*t*} for *T*-contrast or SPM{*F*} for an *F*-contrast). Typically, when using a single univariate statistic one would apply a statistical threshold, or α-level, to make inference over the significance of the test. The standard value used in Psychology is a probability of 95% or an α-level of $p<0.05$. Thus, a difference is considered significant if the chance of that result occurring is less than 5% within the distribution. However, when employing a mass univariate approach, one runs into the issues that surround the making of multiple comparisons. When making multiple comparisons, the probability of a false positive increases as the number of comparisons increases. i.e. if one makes 12 comparisons, when using an α of <0.05, there is a 60% probability that one of those would produce a falsely significant result by chance. In whole-brain fMRI, the mass univariate approach is extended to over 100,000 voxels (in this thesis the smallest number of voxels is 110592).Thus, if using an α-level of $p<0.05$, one would expect 5000 falsely significant voxels, leading to spurious inference about processing in the brain. Correcting for multiple comparisons is therefore essential for meaningful statistical inference in fMRI. One approach to correction is to correct the α-level based on the number of univariate tests performed, giving a new probability value which is used as the threshold. This bonferroni approach divides the α-level by the number of tests to be performed (Dunnett, 1970). Thus, if one performed 100,000 univariate tests (with α at $p<0.05$) the *p*-value would be corrected to $P<0.000005$. Obviously this approach is highly conservative and therefore is also highly likely to result in type II errors. This approach also ignores the spatial information that exists in imaging data. That is, neighbouring voxels in an EPI image are also reflecting neighbouring brain anatomy and correcting the α level, whilst ignoring this information, increases the possibility of false negatives. Another approach is to use Random Field Theory (RFT). Random field theory corrects the *p*-value based on the smoothness of the data, assuming that one activated region should correspond to the FWHM Guassian kernel used to smooth the data during pre-processing. Thus corrections are not made by the number of voxels but the number of resels (or resolution elements) and thus the spatial information within the data is not ignored.

There are alternative approaches which can be used when a specific anatomical hypothesis is not being made; for example, when examining whether activity in any brain area varies with the values from a computational model. One such approach is to correct using the False Discovery Rate (FDR). Unlike Bonferroni correction, which aims to remove all false positives, FDR aims to control the number of false positives (Genovese et al., 2002). FDR uses the probability of false positives amongst the voxels declared positive to correct the threshold. Thus, the α-level is tied to the signal in the data rather than the number of voxels or the smoothness.

If a study has an apriori anatomical hypothesis, it is not always appropriate to correct for multiple comparisons across the whole brain. Instead it is more suitable to make a small volume correction (SVC), where a bonferroni correction is applied only to the number of voxels within a search volume. This approach avoids false negative results, where the null hypothesis is falsely accepted. As such, this is a more appropriate correction for multiple comparisons for the work in this thesis than a whole-brain bonferroni correction. Applying bonferroni correction for the number of voxels in the whole brain could potentially cause falsely negative results within the ACC. In addition, this approach is still considerably more conservative than leaving the α-level uncorrected for multiple comparisons.

2.4.2.4 Group Analysis

As stated above, the aim of the imaging experiments in this thesis is to make inferences about population effects. When one is making inferences at the population level, it is essential to take into account both between and within subject variability. In fMRI, it is possible to do this using a two-stage Random Effects approach (RFX). In a RFX analysis, SPM{*t*} or SPM{*F*} images are created for each subject, either for a particular contrast between conditions or for the main effect of one condition. These images (one per subject per contrast) are then input into another design matrix, where further statistical tests are implemented. This second-level design matrix identifies voxels which are commonly activated for a given contrast across all subjects. This summary statistic approach is seen as highly effective for group analyses of fMRI data (Friston et al., 2005). In this thesis, parametric modulators are often used to examine how activity in some areas of the brain varies statistically with the predictions of a computational model. Thus, these parameters are examined by using a simple contrast between the parameteric regressor and baseline. To display the results of such contrasts, the parameter

estimates for different regressors will be plotted on a histogram. Error bars for these parameters will reflect the standard error of the mean.

2.4.3 Anatomical Localization and Small Volume Correction

An important feature of this thesis is the specificity of the anatomical hypotheses that are being tested. Namely, the studies make specific predictions about information processing in the ACCs and the ACCg. Therefore, a crucial feature of these studies is the ability to distinguish between activity in two neighbouring regions of the cortex. In order to achieve this aim, masks were made of the ACCs and ACCg in each subject. The ACC was defined as the posterior extent of the MCC to the anterior and subgenual extents of the ACC as defined in the introduction (see section 1.3 in chapter one). The posterior extent was defined as a vertical line, 22mm posterior to the Anterior Commisure in MNI space.

The masks were created on the normalized anatomical image of each subject, by hand, using FSLveiw (Smith et al., 2004) and MRIcron (http://www.cabiatl.com/mricro). In order to use these as a SVC mask at the group level, all of the subject specific masks were added together and then used to create an 80% probability mask (i.e. a mask where voxels that were labelled as either ACCg in one mask or ACCs in the other, were only included if they were labelled in 80% of subjects). This threshold was applied to avoid voxels being excluded which fell within the anatomical region of interest. Given the high variability in the anatomy of the ACC, a maximum probability mask would have excluded many voxels which fell within the ACC of the majority of the subjects and therefore the population. These were then used as the search volumes for the SVC. Thus, the SVC approach was specific to the anatomical region to which the hypothesis pertained. In addition to specific hypotheses about the ACC, anatomically specific hypotheses were also made about other areas, based on the results in previous studies. To test hypotheses based on previous literature, masks were created with an 8mm sphere around the reported peak coordinate. Alternatively, when examining a hypothesis pertaining to a previous chapter, a mask of the voxels that were significant in the studies relevant contrast were used. These masks were then used as the search volume for a small volume correction.

In addition to the activity that is examined in the ACCg and the ACCs, activations will also be discussed for which less specific hypotheses are defined (i.e those which survive whole brain

FDR correction). For the localization of these areas, the atlas of Duvernoy and Bourgoiun (1999) was used as an anatomical reference. Results will be displayed on a mean anatomical image in MNI space, created by normalizing each subject's anatomical image and then averaged over subjects. Unless stated otherwise, images will be displayed at the statistical threshold that the result was significant at. For small volume corrected results in the ACC, only voxels that fall within the masks are displayed. For results that are reported as small volume corrected, outside of the ACC, the images are displayed at an uncorrected threshold of $p<0.001$.

2.5 Behavioural Analysis

2.5.1 Statistical Analysis

Each of the studies in this thesis involved the collection of behavioural data. For each of the studies the behaviour was an important index of the subjects understanding of the task. Typically this data was analysed using SPSS. In chapters where repeated measures ANOVAs have been used to analyse behavioural data, Greenhouse-Geisser correction has been applied to the degrees of freedom. This is to compensate for the extent of the violation of the assumption of sphericity. Typically one might use a Mauchly's significance test to determine whether the data violates the assumptions of sphericity. However, Mauchly's is insensitive to such violations when the number of subjects is small. As the number of subjects in each study of this thesis is small, corrections were applied to the degrees of freedom for every repeated measures ANOVA.

2.5.2 Computational Modelling

An important feature of chapters 4 and 6 is the use of computational models to analyse both the behavioural and fMRI data. Both the hyperbolic and the Rescorla-Wagner models (see chapters 4 and 6 for more information on the nature of these models) used in these chapters assume that behaviour is driven by assigning value to available choices (or actions). The actions chosen (or the decisions made) by the subject, reflect the value they have assigned to that choice in comparison to the other available options. It is therefore important to model this action selection process in a way that compares chosen to unchosen actions. One commonly

used approach in Reinforcement Learning Theory (RLT) is the Softmax algorithm (Sutton and Barto, 1981). This algorithm converts values into probabilities i.e. it gives the probability of a chosen action, given its value in comparison to the value of all the available actions. Thus the algorithm takes on the form:

(1)

In this equation the probability, P, of the action, a, actually chosen by the subject on trial, n, is a function of the value of the chosen action over the value of other available actions. The numerator is the exponent of the value, v, of the chosen action, a, divided by a temperature or stochasticity parameter β. The denominator is the sum of the exponents of the values of all the actions each divided by the temperature parameter. In this equation, the coefficient β represents the stochasticity of the behaviour (i.e the sensitivity to the value of each option). As β increases (> 1) all actions tend to equiprobability, with lower values of β amplifying the differences between the values assigned to each choice and increasing the differences between the available choices. The numerator therefore represents the value of the chosen action modulated by how sensitive the subject is to the choice values. The denominator represents the value of both the chosen and unchosen actions divided by how sensitive the subject is to the choice values. Thus, this algorithm assumes that actions become more probable when the value of a chosen action approaches the value of all available actions, i.e when one action has a much greater value than all other available actions. The output of this softmax algorithm is a series of probabilities of the choices made by the subject. This algorithm was used in both of the computational chapters to fit the models to the data.

An important aspect of both the computational models and the softmax procedure is that they contain idiosyncratic free parameters. In this thesis, a maximum log-likelihood approach was used to select the best fitting set of parameters for the behaviour of each subject. This maximum likelihood approach involves varying each of the free parameters within a reasonable range, such that all possible combinations of each of the free parameters are estimated. For each set of parameters the log-likelihood is calculated:

(2) $_{(n)}$

Where the likelihood of each set of parameters (L) is determined by the probability of the performed action, P_a, at trial, *n,* according to the model. The set of probabilities are log-

transformed such that if the model perfectly predicts the data, and $P_a = 1$ on every trial, the Likelihood (L) would be 0. When P_a is less than 1 on any trial, the likelihood of the set of parameters reflecting the underlying decision-making process decreases, and the likelihood (L) assumes negative values. The best fitting parameters were selected using the following algorithm:

(3)

This algorithm identifies the set of parameters for which L was closest to 0 and therefore the parameters which produce the highest probabilities of the choices actually made by the subjects. Where is the parameter set and L is the log-likelihood. This approach has previously been used to fit different types of computational models to different types of decision-making behavior (Brovelli et al., 2008; Pine et al., 2009).

Chapter 3: Prediction errors signal others false beliefs

3.1 Abstract

Interacting successfully in social environments requires the ability to understand that the predictions (or beliefs) of others can be distinct from one's own. False-belief tasks , which test whether an individual is capable of understanding that another's prediction is false, have been used extensively to examine the neural basis of processing other's mental states. However, such studies have not examined brain activity that occurs when another's predictions about the outcomes of decisions are false. In this chapter, an fMRI study is presented which examines the processing of cues which reveal privileged information, only to a scanned subject, about the outcome of a third-person's or computer's decisions. On some trials these cues indicated that the outcome would be unexpected. Subjects performed a false-belief judgement task indicating whether the actual outcome was the same as the predicted outcome. This allowed for testing of the hypotheses that activity in the ACCs would change whenever an outcome was unexpected and that activity in the ACCg would signal the erroneous predictions of a third-person.

3.2 Introduction

One of the aims of this thesis is to examine whether similar computations are performed in the ACCg when processing the outcomes of other's actions, as those that are performed in the ACCs when monitoring the outcomes of one's own actions. As already discussed, neurons in the ACCs have been shown to code prediction error signals in a manner that conforms to the principles of Reinforcement Learning Theory (RLT). Therefore, a particularly pertinent question in the context of this thesis is whether similar signals occur in the ACCg to code for the erroneous predictions of others. To investigate this issue, it is important to use a paradigm in which the predictions of a subject and the predictions of another are distinct from each other at the same point in time. In this study, one of the most well-established paradigms in social cognition research is employed, namely, a false-belief task. However, the typical false-belief design is adapted to the context of reward-related predictions and outcomes. This approach enables the examination of prediction error signals triggered by others' false-beliefs about the outcomes of their decisions.

Understanding that another's belief is false is widely regarded as one of the most sensitive and reliable measures of whether an individual is able to understand the mental states of others (Wimmer and Perner, 1983; Saltmarsh et al., 1995; Perner and Lang, 1999; Saltmarsh and Mitchell, 1999; Hughes et al., 2000; Wellman et al., 2001). As a result, many diagnostic test batteries of disorders of social cognition, such as autism, include different variants of the false-belief task (Hughes et al., 2000). Although aspects of the tests may differ, they share several common properties which allow them to explore whether individuals are capable of representing the mental states of another. The most well-known false-belief task and a useful illustration of their basic principles, is the Sally-Anne paradigm devised by Wimmer and Perner (Wimmer and Perner, 1983). In this task, subjects observe a cartoon with two protagonists, Sally and Anne, who play out one of two scenarios. In both conditions, one character (Sally) places an object (a ball or food item) inside a basket and leaves the room. In the false-belief scenario, another character (Anne) then moves the object to a different location. In the true belief scenario she does not move the object from its original location. After this, Sally returns to the room. The subject is tasked with deciding whether Sally will look for the object in the original location, or where the object now actually resides. If the subject performing the task has a Theory of Mind (ToM), they will realise that Sally's belief about the location of the object is false when Anne has moved the object, but true when the object is in the original location.

The subject should therefore indicate that Sally will always look in the location where she left the object. False-belief trials therefore require the ability to represent the mental states of another, when they are distinct from one's own. As these tasks create situations where the beliefs of two individuals are distinct, they are extremely useful for examining the neural basis of mentalizing abilities.

There is now a considerable body of research which has used neuroimaging methods to examine the neural basis of false-belief processing and mental state reasoning (Frith and Frith, 1999; Frith and Frith, 2003; Frith and Frith, 2006). In these studies, human subjects are required to read stories, cartoons or watch videos and make inferences about the mental states of one or more protagonists (Fletcher et al., 1995; Gallagher et al., 2000; Vogeley et al., 2001; Saxe and Kanwisher, 2003; Saxe and Wexler, 2005; Perner et al., 2006; Saxe and Powell, 2006; Sommer et al., 2007; Hooker et al., 2008; Liu et al., 2009; Miller, 2009; Young and Saxe, 2009; Sommer et al., 2010; Young et al., 2010b; Young et al., 2010a; Zaitchik et al., 2010). In the same manner as the Sally-Anne task described above, subjects are required to indicate whether the beliefs of one of the protagonists are true or false. These studies have identified three interconnected areas which are recruited when processing the mental states of others, namely the temporal poles (TPs), paracingulate cortex and the Posterior Superior Temporal Sulci (pSTS) (Frith and Frith, 2003), i.e. the core circuit which is engaged when processing the mental states of others.

Despite the consistency of activation of the core-circuit across many tasks which engage mentalizing processes, there are limitations to the designs that have been employed. Firstly, in these tasks, subjects are engaged in processing the hypothetical mental state of a protagonist. The subjects are therefore not processing the mental state of a human counterpart with whom they are interacting with in real-time. The second limitation of the studies listed above, is that they have either examined activity that occurs when subjects are indicating their judgement about the true or false nature of the protagonist's belief or alternatively activity aggregated across the whole of a vignette. Thus, although there have been studies which have investigated false-belief processing, none have examined activity that occurs at the exact moment that the true or false nature of the protagonist's belief is made apparent. Another limitation of these studies relates to the specificity of the beliefs to which the tasks pertain. The complexity of the vignettes used in most studies result in multiple true and false beliefs being maintained at the same time. As such, it is not clear in these studies what specific belief the stories are examining. To illustrate this point, in the Sally-Anne task predictions are made about multiple variables including the reward associated with finding the object, the prediction

of the spatial location of the object, the identity of the object etc. More stringent control of the predictions made by a third-person during a false-belief task, may therefore lead to a better understanding of the neural processes underlying mentalizing. The study presented in this chapter examines the processing of false-beliefs, but removes many of the confounds which are found in the false-belief paradigms which have been used to date. It examines the processing of false-beliefs when they pertain to predictions made by others about the outcomes of decisions in real-time. None of the studies listed above have examined whether prediction error signals occur at the point in time when another's prediction (or belief) about the outcome of a decision is false.

Interestingly, false-belief processing and reinforcement learning processes share some similarities. In reinforcement learning, prediction error signals occur when the actual outcome of a decision is discrepant from the predicted outcome (Schultz, 2006). Similarly, in a false-belief task, the prediction of a protagonist is discrepant from the actual outcome they will receive (Wimmer and Perner, 1983). In order to perform the task, the subject must be able to identify this discrepancy between the protagonist's prediction and the actual outcome. For example, Sally's prediction becomes erroneous at the moment in time that Anne moves the object to a new location. Thus, both false-belief processing and reinforcement learning require the ability to identify a discrepancy between a prediction and an outcome. Do prediction error signals occur when a subject knows that another's belief about the outcome of a decision is false?

As stated above, the aim of this thesis is to examine whether comparable information processing occurs in the ACCg and the ACCs. In this study, I examine whether prediction error signals, which are known to occur in the ACCs (Amiez et al., 2005; Matsumoto et al., 2007; Sallet et al., 2007), also occur in the ACCg when the outcomes of another's decisions are unexpected. Subjects monitored the predictions and outcomes of another participant's trials (see table.3.1) (the other participant was actually a confederate and throughout this chapter the term "confederate" will be used to refer to the individual that the scanned subject believed was performing the task with them in real-time). Similar to the Sally-Anne task, the subjects received privileged information, which was not relayed to the confederate. This took the form of cues that the subjects believed were only relayed to them inside the scanner. The cues informed the subject of the outcome of the trial, enabling them to infer whether the predictions of the confederate were true or false. The subjects task was to decide whether this outcome was the same or different from the predicted outcome. Thus, the task was the same of that being performed by subjects performing the Sally-Anne task. On each trial one of two

predictions (positive or negative) could be made by the confederate and there could be one of two actual outcomes (positive or negative). In addition to monitoring the trials of the confederate, the subjects also monitored identical trials performed by a computer. The computer acted as a control agent. This allowed for the examination of trials where there was discrepancy between a predicted and actual outcome, but these factors were not specific to another biological agent's responses (see table.3.1 for design). In this study, activity was examined time-locked to cues which signalled the privileged information only to the scanned subject. This enabled two hypotheses to be tested: firstly, that the ACCs will be activated whenever there is a discrepancy between a predicted and actual outcome of a trial (i.e. on both the confederate and computer trials) and secondly, that the ACCg will be activated exclusively when there is a discrepancy between the predicted and actual outcome on the confederate's trials.

3.3 Methods

3.3.1 Subjects

The subjects were sixteen, healthy right-handed participants (aged between 18 and 30; 9 female), screened for neurological and psychological disorders. One subject (1 female) failed to complete the whole scanning session and was excluded from the analyses. Subjects were paired up with one of two confederate participants, who they believed were a naïve participant. The subjects were not paid for their participation but were offered a picture of their brain as an incentive. The subjects were informed that the other participant performing the task with them (confederate) were being paid £5 for their participation as they were not being scanned.

3.3.2 Training

Subjects were pre-trained in pairs with a confederate one day prior to the scanning session. Training was conducted in two phases.

In the first phase, the subject and the confederate were seated in front of the same monitor, each with their own keypad. They each performed a series of delayed-outcome conditional motor learning trials. During this phase, each trial consisted of an instruction cue (a coloured shape), a trigger cue (three white lines which indicated that a response should be made), the response cue (three white lines with an asterisk over one line indicating which response was made; "missed" was presented at this point if a response was not made within the response window) and finally a feedback cue was presented (a one pound coin indicating a correct response or a one pound coin with a cross through it, indicating an incorrect response, "missed" if a response had not been made at the time of the trigger cue). They were both required to learn the arbitrary stimulus-response associations between three cues and three motor responses by trial and error. They also observed a computer 'learning' associations (a non-biological control, as used in previous studies (Sanfey et al., 2003; Ramnani and Miall, 2004; Apps et al., accepted). Subjects had a 750ms window in which to make a response following the onset of the trigger cue. The instruction cues were colour-coded, such that the subject responded to red shapes, the confederate responded to green shapes and the computer responded to black shapes. However, the form of the shapes was identical and all

three agents learnt the same associations. As the confederates were paired with multiple different subjects throughout the piloting and experimental phases, they were highly over-trained on the associations. However, they were told to make deliberate errors (both responses that were too slow and also incorrect responses) to mimic the learning of a real participant. This first training phase ensured that the subject understood each of the stimulus-response associations and enabled them to observe and monitor the responses of the confederate. Once the subject, the confederate and the computer had made three contiguous, correct responses for each instruction cue, the task was completed.

In the second phase, the subject and the confederate practised the task that would be performed in the scanner on the following day (see task design below). During this phase, the subject was played the sound of the scanners EPI sequence through headphones, inside a mock scanner. The subject observed the confederate being seated in front of a monitor with a response keypad, before they entered the bore of the mock scanner. This practise session lasted 12 minutes and consisted of 90 trials.

3.3.3 Scanning

3.3.3.1 Task Design

		Confederate		***Computer***	
	Prediction	*Positive*	*Negative*	*Positive*	*Negative*
Outcome	*Positive*	*Positive true belief*	*Positive false belief*	*Positive true belief*	*Positive false belief*
	Negative	*Negative false belief*	*Negative true belief*	*Negative false belief*	*Negative true belief*

Table 3.1 Experimental design. A 2x2x2 Factorial design was used. The first factor manipulated Agency (either the confederate or the computer responded on a trial) the second factor manipulated the Prediction (either positive or negative) and the third factor manipulated the actual Outcome (either positive or negative). The colours in the table for each condition match the colour used for each condition in the PSTH plots in the results.

During the second training phase and the scanning session the confederate and the computer continued to make responses on conditional motor learning trials in the same manner as they had during the training session. However, the subject no longer performed trials in the same

manner, instead they performed a false-belief task on the conditional motor learning trials of the other agents (see below). The subjects were informed that they would see the responses of the computer and confederate in real-time from inside the scanner. However, the responses they observed were actually a series of computer controlled responses which occurred in a pre-programmed sequence. To ensure that the subject maintained the belief that the confederate was a naïve participant responding to the visual cues, four missed trials and three incorrect responses were programmed to occur on the confederate's trials. None were included on the computer's trials, to ensure that subject's maintained a sense of biological agency for the confederate and not for the computer. The behaviour of the confederate (the number of errors and missed responses) was based on the responses of subjects in a pilot experiment.

The trials of the confederate and the computer consisted of an instruction cue, a trigger cue and feedback, presented in the centre of the scanned subjects' visual display. All of these trial elements were presented in real-time to the scanned subject inside the bore of the scanner. The same colour-coding of instruction cues was used for the computer and confederate as during training (green for confederate, black for computer). There were 360 trials in total, 180 of which consisted of one of the three instruction cues for which the correct association had been learned during the training session (90 confederate trials, 90 computer trials; 30 trials for each cue for each agent). The subject and confederate were also reminded of the stimulus-response associations they had learned before entering the scanner. However, feedback was pre-determined on these learnt cue trials, such that rewarding outcomes were only delivered on 2/3 of the trials (20 for each shape, for each Agent), even if a correct response was made. Thus, on 1/3 of the learnt cue trials (10 for each shape) a negative outcome was delivered. Subjects were told that a negative outcome did not indicate that a correct response had not been made. A negative outcome indicated that a correct response was unexpectedly not rewarded. Thus, subjects were informed that the correct responses were fixed for the learnt cure trials and that a negative outcome did not indicate that behaviour should be changed on future trials. The confederate was told this information in the presence of the subject. The subject was therefore aware that the confederate would not change their stimulus-response mappings following an unexpectedly negative outcome. The subject was told that the Computer would always make correct responses on learnt cue trials and therefore maintained the same stimulus-response mappings throughout the session.

In addition to the 'learnt' cues, an extra instruction cue was now presented on 50% of the trials (180 trials; 90 computer, 90 confederate trials). For this 'random' cue, the correct response changed randomly between each of the three buttons across trials. As such, it was not possible to learn the correct response for this instruction cue. Thus, the confederate would predict receiving a negative outcome on 2/3 of trials, given that there was a 1 in 3 chance of guessing the correct response on each trial. Unbeknown to the subjects, rewarding outcomes were fixed to be received on only 1/3 of all 'random' cue trials (30 computer trials, 30 confederate trials). The subjects were informed that these positive outcomes were delivered when the confederate or computer had made the correct response. Subjects were informed that although the Confederate and the Computer may be unexpectedly rewarded on these trials when a correct response had been made, the correct response on the next trial was still randomised. The subjects were therefore told that the Confederate was aware that a rewarding outcome was not indicative of the same actions being rewarded or not on future trials. They were also told that the Computer would make a random response on every trial, regardless of the outcome of the previous trial.

In summary, the confederate and computer would receive rewarding outcomes on 2/3 of the learned cue trials and negative outcomes on 2/3 of the random cue trials. Based on probabilities, a positive outcome would be predicted on the learnt cue trials and a negative outcome would be predicted on the random cue trials.

During the scanning session, the subject performed a different task inside the bore of the scanner. During this session, the subject was required to perform a false-belief task. Following the trigger cue, the subject received privileged information that was not displayed to the confederate outside the scanner. The corner of the confederate's monitor was covered during scanning. Before the subject entered the scanner they were shown the monitor that would be used by the confederate and its covered area. By covering the corner of the confederate's screen, additional information could be given to the subject in the corresponding corner of the screen inside scanner. Thus, the subjects were under the belief that they were receiving privileged information that the confederate could not see. On each trial, the subjects received an additional cue (the 'privileged cue') in the corner of the screen, which informed them what the actual outcome of the trial was. At this point in time, the subject knew both the confederate's prediction of the outcome and the actual outcome of the trial. Thus, when there was a discrepancy between the confederate's prediction and the actual outcome of the trial, the subject knew that the confederate was a holding a false-belief. Similarly on all of the

confederate and computer false-belief trials, this event revealed that there was a discrepancy between a prediction and an outcome that was not specific to another biological agent. Following the presentation of this privileged information, a trigger cue appeared in the corner of the screen. At this point the subject was required to indicate whether the confederate held a false-belief. Subjects were instructed to "determine if the outcome is what would be predicted" on the computer trials. Subjects had 750ms to indicate whether the belief was true or false by pressing the first button on the keypad for true and the second button on the keypad for false.

3.3.3.2 Trial Structure (see fig.3.1)

Trials for the subject consisted of an instruction cue (colour-coded shapes for training partner or computer), immediately followed by a trigger cue (instructing a computer or training partner response), followed by a response cue displaying the response of the partner or computer. After a variable delay period a privileged cue was presented (informing the scanned subject what the actual outcome of the trial would be, in the corner of the screen). After a further variable delay the scanned subject trigger appeared (a cue displayed in the corner of the screen, instructing the response from the scanned subject) followed by the scanned subject response (displaying the response of the scanned subject in the corner of the screen). Finally feedback was presented after another variable delay (displaying the outcome of the confederate or computer decision in the centre of the screen).

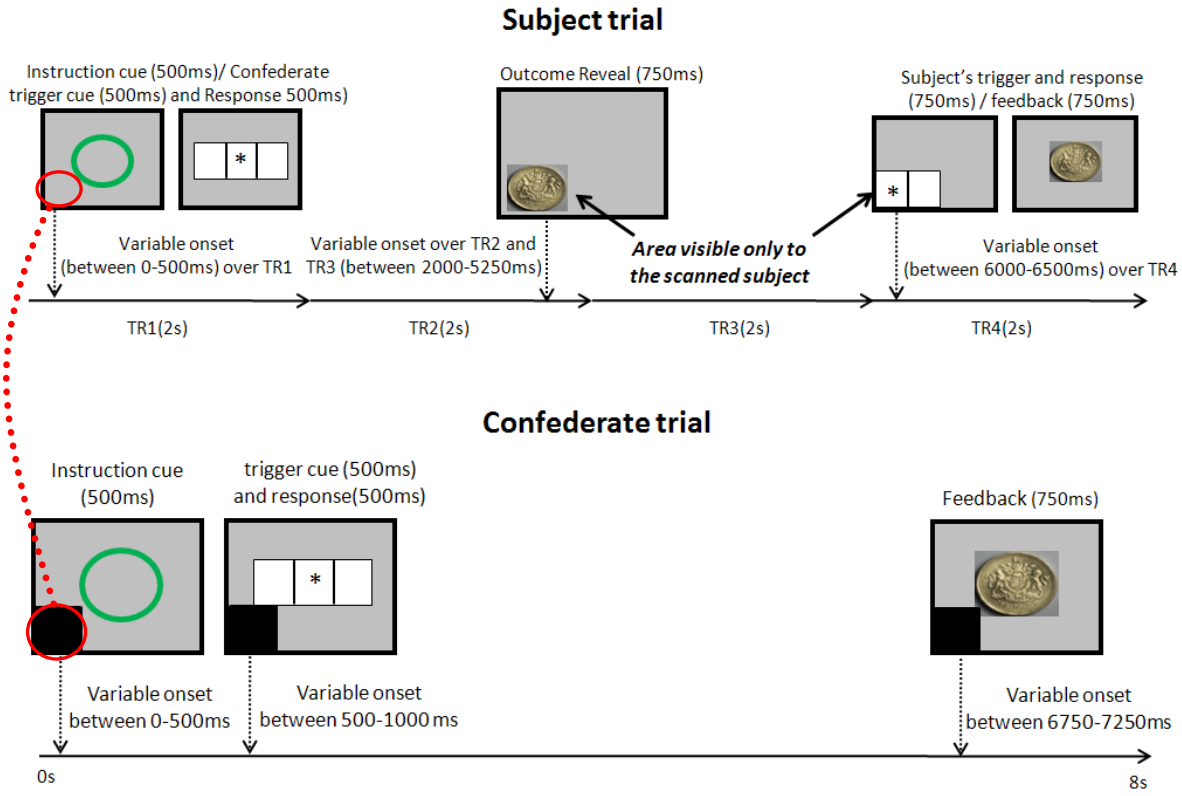

Fig 3.1. Trial Structure for the subject (top) and for the confederate (bottom). The subject saw all of the trial events that were displayed to the confederate, but also observed additional cues that the confederate could not see. The subject monitored the responses of the computer and the confederate. In addition, they received privileged information in the corner of the screen that the confederate could not see outside the scanner (area in black in the corner indicates the area of the screen that could only be seen by the scanned subject). Activity in the ACC time-locked to the outcome reveal event, the point at which predictions could become true or false. The dotted red line and circles highlight the area of the confederate's screen that was covered, which was not covered on the subject's screen.

3.3.3.3 Conditions (see table 3.1)

To investigate activity that occurred at the specific moment in time when new information revealed that the confederate had a false belief, activity time-locked to the 'privileged' cues was examined. A 2x2x2 Factorial design was used. The first factor was the Predicted outcome or 'Prediction' which could positive or negative. The second factor was Outcome, which could be positive or negative and the third factor was Agency which could be either confederate or computer. This created eight different conditions that occurred time-locked to the cues which signalled the privileged information, which were as follows:

1. Confederate positive false belief (the confederate is predicting a negative outcome on the random cue trials, but the actual outcome is positive).

2. Confederate negative false belief (the confederate is predicting a positive outcome on the learned cue trials, but the actual outcome is negative).

3. Confederate positive true belief (the confederate is predicting a positive outcome on the learned cue trials, and the actual outcome is positive).

4. Confederate negative true belief (the confederate is predicting a negative outcome on the random cue trials and the actual outcome is negative).

5. Computer positive false belief (a negative outcome is expected on the computer random cue trials, but the actual outcome is positive).

6. Computer negative false belief (a positive outcome is expected on the computer learned cue trials, but the actual outcome is negative).

7. Computer positive true belief (a positive outcome is expected on the learned cue trials and the actual outcome is positive).

8. Computer negative true belief (a negative outcome is expected on the random cue trials and the actual outcome is negative).

The aim of this investigation was to examine activity occurring when new information revealed another had a false belief and also activity that occurred whenever an outcome was unexpected. To examine these two occurrences, two main contrasts were conducted. The first looked for an interaction between Prediction (positive x negative) and Outcome (positive x negative), independent of the level of Agency. This would identify voxels that showed difference in the BOLD response between true and false belief trials, irrespective of the Agent who made the initial response. In addition to this main interaction, additional contrasts were conducted to look for any main effects. The additional contrasts excluded other possible interactions that could drive the effect that was identified by the main interaction outlined above.

The second contrast looked for an interaction between Prediction (positive x negative), Outcome (positive x negative) and Agency (confederate x computer). The same interaction between Prediction and Outcome for only the computer level of Agency was used as an exclusive mask (P<0.05unc). This ensured that any voxels identified were exclusively responding to the interaction between Prediction and Outcome on the confederate's trials. To determine whether any one condition was driving interaction effects, contrasts were conducted between each confederate false-belief condition and all other conditions. In addition to the these main contrasts, an interaction was carried out between prediction and outcome on the computer's trials, masked by the same interaction on the confederate's trials, to determine if there was any region of the ACC which responded exclusively to computer related responses.

Finally, an additional analysis was performed examining activity time-locked to the instruction cues. At the time of the instruction cues, the subjects would be able to code the Predicted outcome of the trials. To examine activity time-locked to these instruction cues a 2x2 factorial design was used. The first factor was the Prediction (positive or negative) and the second factor was the Agency (confederate or computer).

3.3.3.4 Experimental timing

An important feature of the study was that activity was time-locked specifically to the point in time when privileged information was revealed to the subject. In order to do this, a variable delay was introduced between the instruction cue and the privileged cue. An additional delay was also introduced between the privileged cue and the scanned subject trigger cue. This allowed BOLD activity time-locked to the privileged cue to be isolated, without contaminating effects of either prior or subsequent trial events (Ramnani and Miall, 2003, 2004). Events in each trial took place across four TRs (0–8 s; TR=2s). The interval between scan onset and instruction cue onset was varied over the first TR from trial-to-trial. To optimally sample the cue of interest, the privileged cue, a randomly varying interval between the scan onset and these cues was introduced over the second and third TRs. This achieved an effective temporal sampling resolution much finer than one TR for the conditions of interest. These intervals were uniformly distributed for each condition, ensuring that Evoked Haemodynamic responses (EHRs) time-locked to the privileged cues were sampled evenly across the time period

following each outcome reveal. The scanned subject trigger and feedback cues were randomly jittered over the fourth TR.

3.3.4 Functional imaging and analysis

3.3.4.1 Data acquisition

1470 EPI scans were acquired from each participant. 27 slices were acquired in an interleaved manner, at an oblique angle (≈30˚) to the AC-PC line to decrease the impact of susceptibility artefact in subgenual cortex (Deichmann et al., 2003). A voxel size of 3×3×4 mm (25% slice gap, 0.8 mm) was used; TR=2 s, TE=32, flip angle=80°. The functional sequence lasted 49 minutes. High resolution T1-weighted structural images were also acquired at a resolution of 1×1×1 mm using an MPRAGE sequence.

3.3.4.2 Image preprocessing

Scans were pre-processed using SPM5 (www.fil.ion.ucl.ac.uk/spm) by spatial realignment to the first scan, normalization to the ICBM EPI template using both linear affine transformations and non-linear transformations (Friston et al., 1995a). Lastly, a Gaussian kernel of 8 mm was applied to spatially smooth the images in order to conform to the Gaussian assumptions of the GLM implemented in SPM5.

3.3.4.3 Event definition and modelling

Nine separate event types were modelled in the analysis. Each of the eight conditions time-locked to the privileged cue, were modelled as a separate event type. The instruction cues, trigger cues, feedback cues, and the privileged cues from trials which were either missed or incorrect responses, were modelled as one regressor. Trials were classed as missed if the response was too early (before the scanned subject's trigger cue) or too late (a reaction time > 750 ms). Each event type was used to construct a series of regressors by convolving the event timings with a Fourier set (see below) of five harmonic functions (two sine, two cosine, one envelope function with a Hanning window of 32s). The residual effects of head motion were modelled in the analysis by including the six parameters of head motion acquired from the

realignment stage of the preprocessing as covariates of no interest. Prior to the study, a set of planned experimental timings were carefully checked so that they resulted in an estimable GLM in which the statistical independence of the eight event types was preserved.

In this study, a Fourier basis set was used to examine EHRs following events rather than the canonical HRF which was described in chapter two . This circumvented the potentially false assumptions that are made when using the canonical HRF to model the BOLD response. The canonical HRF assumes that the BOLD response peaks at ~6s and undershoots at ~12s. However, the timecourse and form of the BOLD response may not always fit this response profile across the whole brain (Henson and Rugg, 2001; Henson et al., 2002). The stimulus evoked BOLD response can vary across the brain and across subjects (Handwerker et al., 2004), with deviations from this profile identified in the ACC (Lau et al., 2006), which is particularly pertinent for this study. Use of the canonical HRF can therefore result in type II errors in which real differences in the BOLD response between conditions are not identified because the response profile of that brain area differs from that of the basis set used. An alternative approach is to convolve a set of Fourier functions to the event onsets. This approach is more flexible in the form and the timecourse of the BOLD response following an event. This approach has been validated and used successfully in the past to examine event-related BOLD responses (Ramnani and Miall, 2003, 2004; Balsters and Ramnani, 2008).

A secondary analysis was conducted to examine activity time-locked to the instruction cues. For this, a separate design matrix was created in which the onsets of instruction cues were separated into a 2x2 design. The first factor was the Prediction (positive or negative) and the second was Agency (computer x confederate). The instruction cue onsets were fitted with a Fourier set of five harmonic functions. The onsets of the different conditions at the time of the privileged cue were collapsed into one regressor which modelled all responses to that event. The rest of the design matrix remained the same as outlined above.

3.3.4.4 First-level analysis

The GLMs were estimated in SPM5 (Friston et al., 1995b). SPM{*t*} contrast images were computed at the first-level, one image per basis function. In the primary analysis, 45 SPM{*t*}

images were created at the first-level to be used in the second level. In the secondary, 25 SPM{*t*} images were taken to the second-level.

3.3.4.5 Random effects group analysis

A random effects analysis (Full-Factorial ANOVA) was applied to determine voxels significantly different at the group level. SPM{*t*} images from all subjects at the first-level were grouped into two factors, basis function and condition. *F*-contrasts were conducted across the Fourier basis functions to look for significant interaction effects (see 'conditions' above). To apply correction for multiple comparisons, we used 80% probability anatomical masks of the ACCg and ACCs. To create each mask, subject-specific masks of the ACCg and ACCs were constructed in FSL using the anatomical criteria outlined in chapter two. The ACCs mask was to correct for multiple comparisons in the contrasts examining all unexpected events and the ACCg mask was used in the contrast examining the confederate's false beliefs.

3.4 Results

3.4.1 Behavioural Results

Subjects performed a false-belief task on the predictions of a confederate and also trials performed by a computer. On each trial they were presented with a cue which signalled the actual outcome of the trial. This cue was privileged and not presented to the confederate. At this moment in time the subject was able to determine whether the actual Outcome of the trial was discrepant from the Predicted outcome. On each trial, following the presentation of this cue, the subjects were required to indicate whether the Predicted outcome was the same as the actual Outcome. Subjects performed the task at a high level of accuracy (mean of 92.9% of 355 trials performed correctly; mean 25.2 trials incorrect or missed SD±13.54). Thus, subjects were able to correctly understand both the predicted and actual outcomes of trials. To examine the subjects' performance of the task in both the confederate and computer conditions the number of correct trials in each condition were converted into an overall percentage for the confederate and computer trials. To test for any significant difference in task performance between the confederate and computer trials, a repeated measures *t*-test was conducted. No significant difference in task accuracy was found between the computer and confederate conditions ($t(14) = 0.174$, P=0.865). Therefore, subjects were able to perform the task at the same level of accuracy regardless of whether it was a computer or confederate trial.

3.4.2 Imaging results

3.4.2.1 Instruction cue-related activity

To examine prediction related activity time-locked to the instruction cues, three F-contrasts were conducted. The first contrast looked for a main effect of Prediction (Positive <> Negative), the second looked for a main effect of Agency (computer <> confederate) and the third looked for an interaction between Prediction and Agency. The results showed that there was no main effect of Prediction, or Agent and also no interaction between these factors in the whole-brain analysis (FDR corrected). In addition, there were no effects of Prediction, Agent or an interaction between these factors in the ACCs or ACCg when using a small volume correction. Thus, no significant effects were found time-locked to instruction cues.

3.4.2.2 Privileged cue-related activity

The main aim of this experiment was to investigate activity time-locked to events that signalled an error in another's prediction. A 2x2x2 factorial design was used to examine activity time-locked to the cue which signalled the privileged information only to the scanned subject. The first factor was the predicted outcome of the trial (Prediction), which could be positive or negative; the second factor was the actual outcome of the trial (Outcome), which could be positive or negative and the third factor was the respondent on the trial (Agent) which could be either the confederate or the computer.

The experimental design contained four conditions in which the actual outcome was discrepant from the predicted outcome. This included two conditions in which the confederate's prediction of the outcome would be false and two conditions in which an unexpected outcome occurred on the computer's trials (see fig.3.1). Two hypotheses were tested: firstly, that the ACCs would respond on every trial where the actual Outcome was different from the Prediction, regardless of whether it was a computer or confederate trial. Secondly, the ACCg would respond exclusively when the confederate's prediction was false. To stringently test the anatomical hypotheses, masks of the ACCg and ACCs were used as a small volume correction for multiple comparisons. These masks ensured that any activated voxel at the group level would be within the ACCg or ACCs of 80% of the subject.

3.4.2.2.1 Prediction error activity

To test the first hypothesis, I looked for an interaction between Prediction and Outcome independent of whether the trial was that of the computer or the confederate. This contrast identified voxels which showed a significant effect to any unexpected Outcome, regardless of the Agent performing the trial. A significant effect was found in the ACCs (putatively within the Rostral Cingulate Zone (RCZ) in midcingulate area 24c'). Examination of peristimulus time histograms (PSTH), of data from the peak voxel, revealed that this effect was being driven by a significant response to both the confederate and computer positive false belief trials (see fig.3.2). To test whether the interaction effect in this peak voxel was being driven by the response to these two conditions, additional contrasts were conducted (Table.3.2). These contrasts revealed significant differences between the computer and confederate positive false belief conditions, where the outcome was unexpectedly positive, and all other conditions. Thus, the interaction was driven by responses to both the confederate and computer positive

false belief conditions. Importantly, there was no significant difference between these two conditions, even at a much lower threshold (F(5,560) = 0.52; Z = 0.1; *p*=0.532unc). This indicates that the ACCs made a response on all trials where there was an unexpectedly positive outcome.

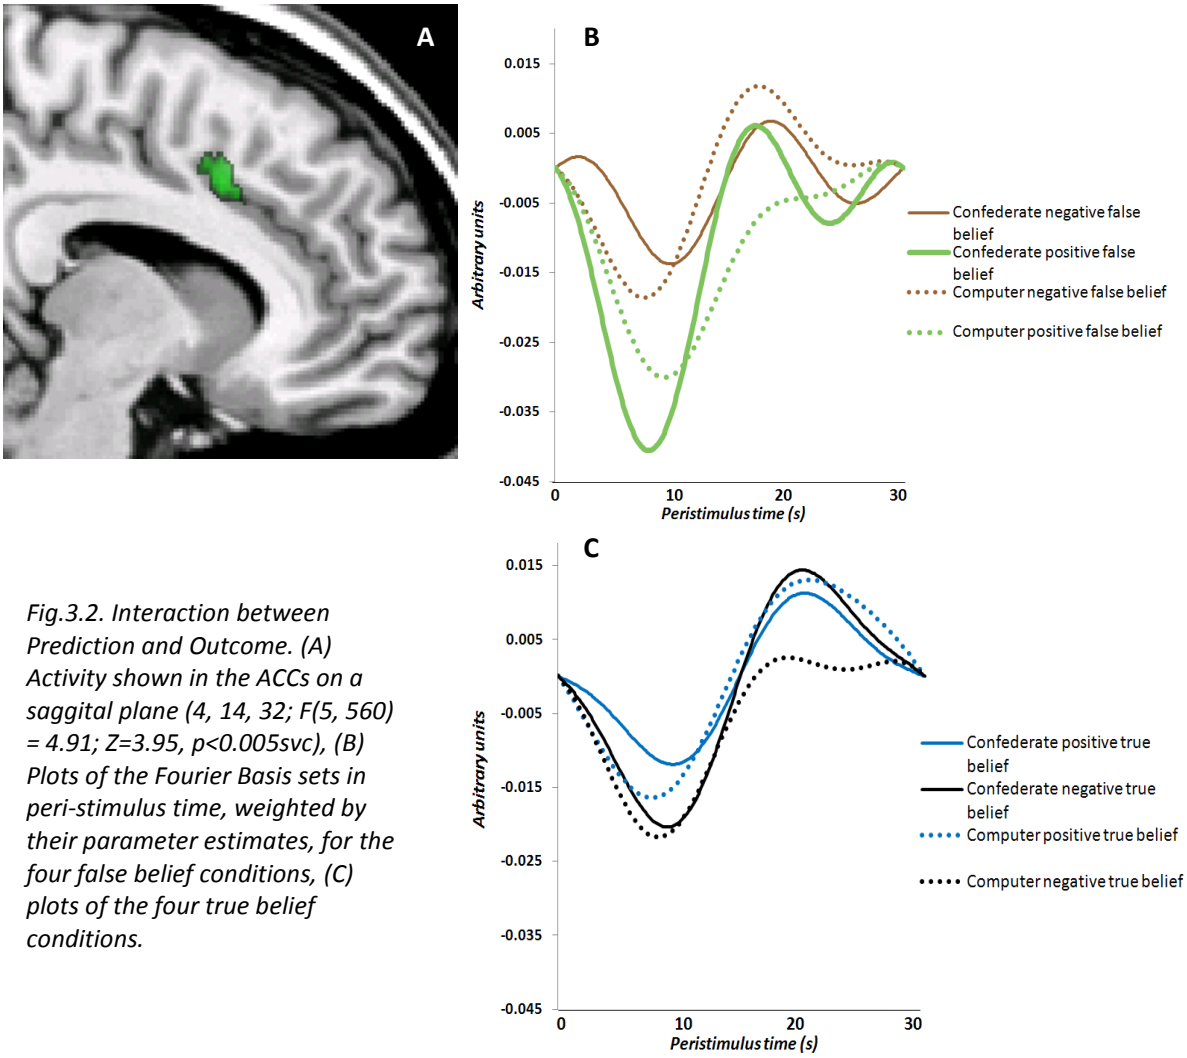

Fig.3.2. Interaction between Prediction and Outcome. (A) Activity shown in the ACCs on a saggital plane (4, 14, 32; F(5, 560) = 4.91; Z=3.95, p<0.005svc), (B) Plots of the Fourier Basis sets in peri-stimulus time, weighted by their parameter estimates, for the four false belief conditions, (C) plots of the four true belief conditions.

Contrast	Z	P
(Confederate and Computer positive true belief) <> (Confederate and Computer positive false belief)	3.48	P<0.005
(Confederate and Computer negative true belief) <> (Confederate and Computer positive false belief)	3.48	P<0.005
Confederate and Computer negative false belief) <> (Confederate and Computer positive false belief)	3.22	P<0.005
Confederate positive false belief <> Computer positive false belief	1.73	P>0.01

Table 3.2. Contrasts conducted to examine the interaction effect identified in the ACCs .This table shows that the peak voxel in the ACCs cluster shows a significant difference between the two types of unexpectedly positive outcome (the confederate and computer positive false belief conditions) and the other conditions. The top three results were small volume corrected for multiple comparisons, the bottom result is reported uncorrected. The degrees of Freedom for all comparisons were F(5,560).

3.4.2.2.2 Confederate false-belief

To test the second hypothesis, I looked for an interaction between Prediction, Outcome and Agency. To ensure that voxels which showed an interaction effect between Prediction and Outcome on the computer trials did not contribute to this effect they were "masked out" (i.e an exclusive mask of the uncorrected (p<0.05) interaction between Prediction x Outcome for only the level of computer was used). This ensured that only voxels which showed the interaction exclusively for the confederates' conditions would be identified. A significant effect was found in the ACCg (0,8,28, F(5, 560) = 5.2; Z=3.69; p<0.05svc; see fig.3.3) putatively in midcingulate area 24a'/24b'. Notably, this area did not overlap with the cluster that was activated by any unexpectedly positive outcome reported above (see fig.3.4). Examination of the PSTH revealed this effect was driven by a response only to the confederate positive false belief condition. To test statistically whether this interaction effect was being driven by this response, contrasts were run between this positive false belief condition and every other main effect. All seven of these contrasts showed a significant effect (F(5,560); p<0.05 svc) in the peak voxel identified in the interaction above. These results indicate that a portion of the ACCg responds whenever a subject knows another will unexpectedly receive a positive outcome.

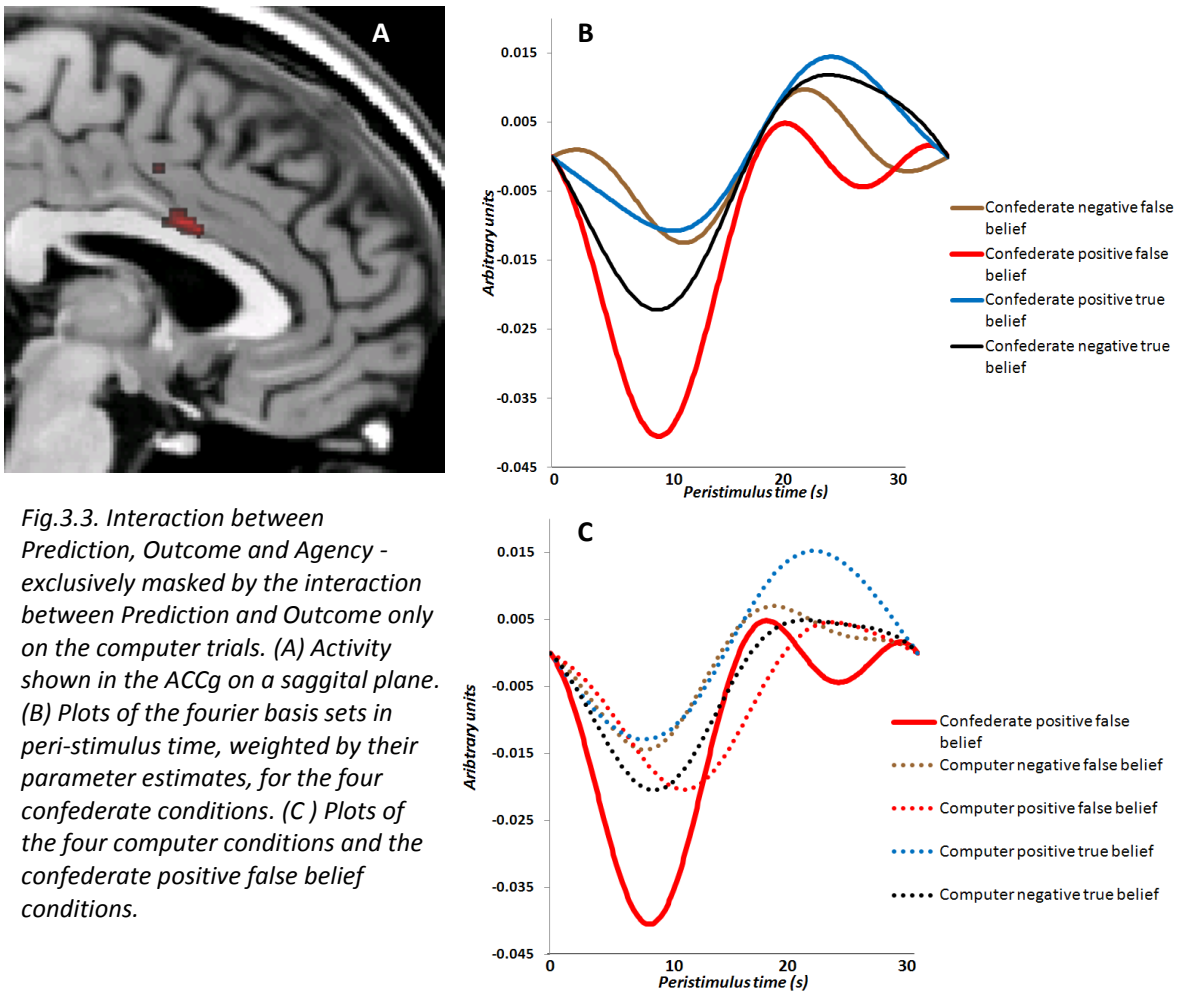

Fig.3.3. Interaction between Prediction, Outcome and Agency - exclusively masked by the interaction between Prediction and Outcome only on the computer trials. (A) Activity shown in the ACCg on a saggital plane. (B) Plots of the fourier basis sets in peri-stimulus time, weighted by their parameter estimates, for the four confederate conditions. (C) Plots of the four computer conditions and the confederate positive false belief conditions.

Essential to the interpretation of these results is the anatomical specificity of the activity in the sulcus and in the gyrus. It was of particular importance to ensure that the effects specific to the confederate prediction errors were within the gyrus and not the sulcus. To test this possibility that that a portion of the ACCs signalled the false-beliefs of the confederate, I used the 80% ACCs probability mask on the confederate interaction contrast, i.e. used a small volume correction of voxels in the sulcus on the contrast examining confederate prediction error signals, where the hypothesis related to the gyrus. No significant voxels were present. To examine the alternative possibility, that voxels within the ACCg signal any unexpected positive outcome, I used the ACCg mask on the interaction between Prediction and Outcome. Three voxels showed a significant effect. This potentially confounds the results described above, as there are voxels within the ACCg which signal any unexpected outcome. However, it should be noted that these 3 voxels were not spatially contiguous, lay at the extremities of the mask (i.e each was at the most lateral extent of the mask closest to the sulcus) and were at the border between the ACCg and the ACCs masks. Thus, given the lack of spatial resolution of smoothed EPI data and the lack of an anatomical boundary between the sulcus and the gyrus by which to define the masks, it is possible that these voxels did not fall within the ACCg.

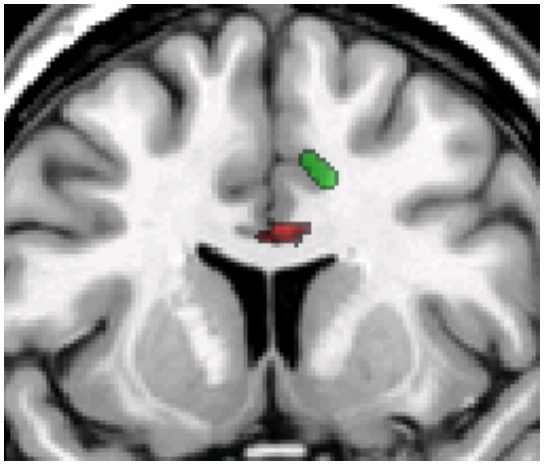

Fig.3.4 Results of the prediction error and false belief contrasts shown on a coronal place (Y = 12). A significant response occurred in the ACCg (shown in red) for the confederate positive false belief conditions. A significant response occurred in the ACCs (shown in green) to both the computer and confederate positive false belief trials.

In summary, these findings show that parts of the ACC have the common property of sensitivity to unexpectedly positive outcomes of decisions. However, responses of the ACCs and ACCg differ in terms of whose errors are being processed. Whereas the ACCs responds to any unexpectedly positive outcome, the ACCg responds exclusively when the outcomes of another's decision is unexpectedly positive.

3.5 Discussion

This study tested two hypotheses about the processing of unexpected outcomes and false beliefs in the ACC. The first hypothesis was that the ACCs would be engaged at the particular point in time that new information revealed that the actual outcome of a decision would be different from the predicted outcome. The second hypothesis was that the ACCg would be activated at the specific point in time that this information revealed that another's prediction was erroneous. The results show that both the ACCg and the ACCs respond to unexpectedly positive outcomes. In line with the hypotheses, the ACCs is sensitive to any unexpectedly positive outcome, however the ACCg is sensitive specifically to the unexpectedly positive outcomes of another's decisions.

3.5.1 Prediction errors and the ACC

This study showed activation in the ACC when predictions about the outcomes of decisions were erroneous. There is a considerable body of evidence which shows that the ACC is involved in processing the outcomes of decisions and more specifically in signalling when they are different from expectations (Carter et al., 1998; Walton et al., 2004; Frank et al., 2005; Behrens et al., 2007; Walton and Mars, 2007; Rushworth and Behrens, 2008; Kennerley et al., 2009; Kennerley and Wallis, 2009a). Single-unit recording studies have shown that neurons in the ACC show an increased firing rate as the probability of reward increases (Shidara and Richmond, 2002; Sallet et al., 2007; Kennerley et al., 2009; Kennerley and Wallis, 2009a; Hayden and Platt, 2010) and also increase their firing rate in line with the magnitude of the prediction error responses (Amiez et al., 2005; Matsumoto et al., 2007; Sallet et al., 2007). This would suggest that the ACC codes reward predictions and errors in the manner proposed by Reinforcement Learning Theory (RLT). That is, there are neurons which code predictions about the outcomes of choices during choice selection and also neurons in the same region which signal when the outcome of the decision reveals that predictions were erroneous. One might argue that such prediction error signals, found in a small sample of neurons in these studies, may not be informative as to the function of the ACC as a whole. However, there is also converging evidence from studies investigating the behavioural effects of lesions in non-human primates, as well as electrophysiological and neuroimaging studies in humans. These studies suggest that the error signals in the ACC may be involved in learning and guiding

behaviour. For instance, lesions to the ACCs in monkeys result in an inability to adapt behaviour appropriately in a task requiring reinforcement learning processes (Kennerley et al., 2006). In humans, the ACC is believed to be the dipole source of the error-related negativity (Frank et al., 2005; Holroyd and Coles, 2008; Holroyd et al., 2009), which occurs in many situations were an error is experienced, including instances when a positive outcome is received when a negative outcome was predicted (Frank et al., 2005). fMRI studies in humans also show an analogous BOLD response that occurs whenever an outcome of a decision reveals a prediction was erroneous (Carter et al., 1998; Holroyd et al., 2004; Jessup et al., 2010; Nee et al., 2011). As such, the signal identified in the ACCs in this study, closely matches many of the findings in the literature. However, in this study the signal occurred whenever there was an unexpectedly positive outcome on computer or third-person's trials. This is distinct from the findings of previous studies, which have shown such signals to occur when the unexpected outcome is that of the subjects' own decisions. This study also provides evidence for the hypothesis that is outlined in chapter one. Similar information processing is performed in the ACCs and the ACCg, but the ACCg is processing this information about the decisions of another.

The view that the ACCs and the ACCg perform similar reward-related computations is supported by anatomical evidence. Both areas receive direct projections from dopamine neurons within the VTA (Williams and Goldman-Rakic, 1998) and also projections from the striatum via the pathway through the ventral pallidum and the thalamus (Yeterian and Pandya, 1991). Importantly, the VTA and the striatum have been shown to signal reward prediction errors (Fiorillo et al., 2003; McClure et al., 2003; O'Doherty et al., 2003; Tobler et al., 2005; Seymour et al., 2007; Brovelli et al., 2008; D'Ardenne et al., 2008). Such a 'connectional fingerprint' implicates the ACC in processing information related to the prediction error signals in the ventral striatum and VTA. In addition, both the ACCs and ACCg share connections to areas implicated in decision-making and reward processing, portions of the intraparietal sulcus (Pandya et al., 1981; Vogt and Pandya, 1987; Petrides and Pandya, 2007), the hippocampal formation (Vogt and Pandya, 1987) and parts of the orbitofrontal cortex (Pandya et al., 1981; Carmichael and Price, 1995). However, despite many common connections the ACCg and the ACCs also exhibit several distinct connections. Firstly, the ACCs shows strong projections to the motor system (Dum and Strick, 1991; Picard and Strick, 1996; Takada et al., 2001) which the ACCg does not. In contrast the ACCg shows distinct connections to medial parts of the superior frontal gyrus (bordering areas 8, 9 and 32'), portions of the pSTS, at the tempero-parietal junction (TPJ) and the TPs (i.e. the core-circuit engaged when mentalizing) which are not found in the ACCs (Pandya et al., 1981; Vogt and Pandya, 1987; Seltzer and Pandya, 1989; Petrides

and Pandya, 2006). This would suggest that the ACCg exchanges information with a circuit which processes the mental states of others, but also accesses reward-related information, in the same manner that the ACCs has access to reward-related information. Thus, anatomical evidence supports the notion that the ACCs and the ACCg signal the discrepancy between a prediction and an actual outcome. However, they process this information for one's own predictions, or the predictions of another, respectively.

As stated above, the results of this study conform to some of the predictions of the RLT of decision-making. More specifically there are signals which code for discrepancies between predictions and outcomes. Such signals are also hypothesised in many of the computational models which are used to examine reinforcement learning behaviours. This study did not use a modelling approach. However, it did examine signals that would be predicted by such models and therefore adds to the growing evidence that the neural activity which underlies decision-making can be explained by computational models (Frank et al., 2005; Doya, 2008; Rushworth and Behrens, 2008). Recently, a small number of studies have shown that signals predicted by RLT occur when subjects were making decisions during social interactions. A number of these studies have shown that the ACCg is engaged when there are discrepancies between the predictions of a subject and the actual outcomes of decisions in social situations (Behrens et al., 2008; Hampton et al., 2008; Baumgartner et al., 2009). However, this study showed that the processing in the ACCg may be specifically to compute the discrepancy between another's prediction and the actual outcome of their choices. This research therefore highlights the utility of RLT in explaining signals that occur during learning in both social and non-social situations.

3.5.2 False-belief and autism

Typically, studies that examine the processing of false beliefs report activation in three interconnected areas, namely the paracingulate cortex, the pSTS and the TPs (Fletcher et al., 1995; Gallagher et al., 2000; Frith and Frith, 2003; Saxe and Kanwisher, 2003; Grezes et al., 2004; Perner et al., 2006; Sommer et al., 2007; Aichhorn et al., 2009). To date, only one of these studies has shown activation in the ACC when processing other's false beliefs (Sommer et al., 2007). The discrepancy between the findings of this study and those of previous studies can be attributed to two unique features of the design.

Firstly, none of these studies have investigated brain activity that occurs at the exact point in time when new information reveals that another's belief is false. Instead they have investigated activity aggregated across entire trials, blocks of trials, or examined activity time-locked to retrospective false-belief judgements. Thus, they could not attribute activity specifically to the event which signalled a false-belief and lacked sensitivity for identifying activity in the ACC time-locked to this event. In this study, activity was time-locked to the particular event that allowed a subject to infer the mental state of another. Thus, unlike previous studies, activity was examined at the exact moment in time when the discrepancy between the beliefs of two individuals was revealed.

Secondly, in most false-belief tasks the subject does not reason about the mental states of another individual whilst interacting with them in real-time (Fletcher et al., 1995; Gallagher et al., 2000; Saxe and Kanwisher, 2003; Perner et al., 2006). Instead, the subject is reasoning about the hypothetical mental state of a protagonist. In this study, the subject was making inferences about specific predictions made by another person whilst monitoring their behaviour in real-time. Thus, when privileged information revealed that the outcome of another's decision was to be unexpected, it indicated that their current mental state or prediction was false. Such differences between other false-belief paradigms and the one used in this study, may therefore explain the activation within the ACCg identified in this study that has not been found by others.

False-belief paradigms are one of the most commonly used tools to diagnose individuals with disordered ToM processing and investigate the age at which individuals acquire ToM abilities (Saltmarsh et al., 1995; Perner and Lang, 1999; Saltmarsh and Mitchell, 1999; Wellman et al., 2001; Onishi and Bailargeon, 2005; Southgate et al., 2007; Surian et al., 2007). Most test batteries used to diagnose individuals with Autism Spectrum Disorders (ASDs) contain a false-belief test (Hughes et al., 2000). A failure to understand another's false belief indicates that an individual is not able to represent the mental state of others. Interestingly, it has been found that patients with ASDs have disturbed cytoarchitecture in areas 24a' and 24b' which lie predominantly in the ACCg (Palmen et al., 2004). Post-mortem analyses have revealed that ASD patients have global decreases in cell density and cell size, with specific decreases in the quantity of Von Economo Neurons (VEN) compared to a matched control group (Nimchinsky et al., 1999; Simms et al., 2009). Thus, failures of ASD patients to pass false-belief tasks may partially be a result of disturbed neurobiological properties in the ACCg. Unfortunately, however, the limitations of the false-belief studies conducted to date have left discussions of

the role of the ACCg in performing such tasks largely absent. As a result, this area has also not been considered within the most prominent theories of the social cognitive deficits of ASD patients (Baron-Cohen et al., 2000; Williams et al., 2001). The results reported here are therefore particularly significant for theories of disturbances in social cognitive processes in the Autistic spectrum.

3.5.3 Caveats and Limitations

An important aspect of the results was that no significant effects were identified within the ACC for either the confederate or the computer negative false belief trials, i.e. the ACC did not show a response to unexpectedly negative outcomes. At first this seems somewhat surprising, as there are neurons in the ACC which code both positive and negative prediction errors (Amiez et al., 2005; Matsumoto et al., 2007; Sallet et al., 2007). However, this may be accounted for by a distinction between the decision-outcome contingencies on the trials where a positive outcome was predicted and those on which a negative outcome was predicted. Specifically, on the random trials where the stimulus-response mapping could not be learned, a negative outcome indicated an incorrect response and a positive outcome indicated a correct response. These outcomes were therefore contingent on the responses made. However, for the learnt trials, a negative outcome occurred on 1/3 of the trials regardless of whether the correct response had been made. As such, on the negative false belief trials, there was no direct mapping between a correct response and a positive outcome, whereas on the positive false belief trials such a response-outcome contingency was always present. It is well documented that in the VTA, prediction error signals occur only when there is a contingency between a choice and an outcome (Fiorillo et al., 2003; Schultz, 2006). Similarly, fMRI studies have shown that ACC activity is increased when there is a contingency between actions and their consequences (Walton et al., 2004; Jocham et al., 2009). In addition, lesions to the ACCs impair reinforcement learning processes and not the ability to detect erroneous responses per se. (Kennerley et al., 2006; Rudebeck et al., 2008). In this study, the lack of contingency between the choice made and the outcome on the negative false belief trials, meant that no new information was learnt at the time of these outcomes. The lack of contingency between a response and an outcome may therefore explain the absence of a response in the ACC on those trials.

An additional caveat must be noted in relation to the interpretation of the results offered above. Previously I suggested that the signals identified in the ACC are akin to the prediction error signals that are proposed by RLT. However, in this study, the confederate was not explicitly engaged in a learning task, although, learning may have been occurring on the trials where there was contingency between a response and an outcome. Thus, although there were discrepancies between predictions and outcomes, these were only one-shot discrepancies and the value of an action was not explicitly being updated. This deviates from the strictest interpretation of RLT, which suggests that prediction error signals are used for learning and guiding future choice behaviour (Sutton and Barto, 1981; Schultz, 2006). However, in the next chapter I provide evidence to suggest that the identified signals in the ACCg, conform to the computational principles of RLT.

The final caveat that should be noted relates to the absence of any significant effects time-locked to the instruction cues. One would typically expect signals to occur in a widespread network of brain areas for the prediction of a rewarding outcome (Knutson and Cooper, 2005; Schultz, 2006; Hampton and O'Doherty, 2007) and, in particular, one would expect activation in the ACC (Rogers et al., 2004; Hampton and O'Doherty, 2007). However, it is highly likely that the absence of any effects can be attributed to the lack of power afforded by the experimental design for examining instruction cue-related activity. In this experiment, the stimulus timings employed were aimed specifically at maximising sampling of the BOLD response following the privileged information. Therefore, to maximise sampling of the privileged cues, the jitter for the instruction cues was limited to just 750ms, over one TR. As a result, the BOLD response following the instruction cues was both unevenly and poorly sampled. The absence of any statistically significant effects may therefore have been a result of the inability to detect such signals, rather than the absence of changes in neural activity coding reward predictions. However, such an approach was necessary to maximise the sampling of the cue of interest (the 'privileged cue') and was therefore a worthwhile sacrifice in order to test hypotheses about prediction error signals in the ACC.

3.5.4 Summary

In summary, this study examined activity in the ACC during a decision-making based false-belief paradigm. Activity was examined at the point in time when privileged information about the outcomes of decisions was revealed to a scanned subject. The results showed that the ACCs and the ACCg were engaged when the outcome of a trial was unexpectedly positive.

However the ACCg responded exclusively when the unexpectedly positive outcomes occurred as the result of another's decision. These results suggest that prediction error signals, as hypothesised by RLT, occur in the ACCs and the ACCg. However, prediction error signals in the ACCg code for other's false beliefs.

Chapter 4: ACCg prediction error signals when instructing others

4.1 Abstract

In many species, learning often occurs via the transmission of information from individuals with expertise, to novices. Successfully instructing another conspecific requires an expert teacher to identify how erroneous the knowledge of a novice student is and instruct them as to the accuracy of their knowledge. In the previous chapter, prediction error signals were identified in the ACCg when a subject received information that informed them another's prediction was erroneous. Do such signals code for the erroneous predictions of a student in the brain of a teacher? In this chapter, an fMRI study is reported which examines the processing in the brain of a teacher when they are instructing a novice student. Subjects learnt 10 arbitrary visuomotor associations between instruction cues and actions in a training session. They were then scanned whilst acting as a teacher providing error feedback to a novice student who was learning the same associations. Activity was examined in the brain of the teacher at the moment in time when they saw the response of the student. A Rescorla-Wagner algorithm was fitted to the behaviour of the student. The prediction error parameter from this model was then fitted to the responses of the student. This allowed for testing of the hypothesis that the ACCg in the brain of the teacher would signal the erroneous predictions of the student.

4.2 Introduction

An aim of this thesis is to investigate the contribution of the ACCg to processing the decision-making of others. The previous chapter showed that activation of the ACCg occurs whenever privileged information reveals that another's prediction of a negative outcome was false. The result was interpreted within the framework of Reinforcement Learning Theory (RLT). Specifically, it was suggested that the ACCg signals coded for the discrepancy between the prediction of the other person and the actual outcome known only by the subject at that time. The caveat to that interpretation was that the prediction and outcome on one trial were not dependent on the prediction and outcome of the previous trial. As such, the subject was engaged in a "one-shot" learning process. This does not fit with strictest interpretation of RLT which suggests that the function of error signals is to correct erroneous predictions, in order to optimise future decisions. Thus, the one-shot learning design in the previous chapter does not allow for the principles of RLT to be stringently tested. This chapter will test whether activity in the ACCg conforms to these principles when subjects compute the discrepancy between another's prediction and their outcome. One form of social interaction which may require such processing is during instruction, when a subject is required to monitor the learning of another and inform them of their accuracy of their behaviour.

In many species, learning often occurs via the transmission of information from individuals with expertise, to novices (Franks and Richardson, 2006; Thornton and McAuliffe, 2006; Hoppitt et al., 2008). In some situations the behaviour of one agent is aimed explicitly at facilitating the learning of another, i.e. one agent teaches another. There is a growing body of neuroimaging research examining the neural basis of social learning. Such studies have examined activity in the brain of the learner either during observational learning or imitation, or when recalling information learned from others (Buccino et al., 2004; Leslie et al., 2004; Iacoboni et al., 2005; Jackson et al., 2006; Reithler et al., 2007; Monfardini et al., 2008; Shane et al., 2008; Gazzola and Keysers, 2009; Burke et al., 2010; Kang et al., 2010). To date, there is very little research investigating the neural mechanisms that underpin the teaching process. The learning which occurs in studies which investigate social learning processes is passive, with the behaviour of the expert not aimed explicitly at correcting errors committed by a learner. Thus, there is very little research into the neural and computational mechanisms which underpin the behaviour of someone who is engaged explicitly in facilitating the learning of another.

The aim of any teaching process is for the learner to acquire the same expertise as those instructing them (Caro and Hauser, 1992). In order to accurately facilitate learning in another, a teacher must monitor the responses of a student, determine the extent to which their response is different from the correct response, and then provide them with the appropriate feedback. The feedback enables the student to positively or negatively reinforce the chosen behaviours, so that in the future they perform the same responses as would be performed by the teacher. Thus, students learn from a teacher in a manner that conforms to the principles of RLT. I suggest that instructing a student, therefore requires the teacher to model the learning of a student. The teacher must inform the student when their predictions are discrepant from the outcome known by the teacher. However, no previous study has examined whether teachers model the reinforcement learning processes of their students.

The traditional account of reinforcement learning processes suggested that prediction error signals occurr when predictions about rewarding outcome are erroneous. However, there is considerable evidence that there are many different types of prediction error signals that conform to the same computational principles. These error signals pertain to different forms of prediction and may be coded for in separate neural circuits (Dayan and Daw, 2008; Rushworth et al., 2009). Neurophysiological investigations have shown that prediction error signals occur in the Ventral Tegmental Area (VTA)(Mirenowicz and Schultz, 1996; Montague et al., 1996; Schultz et al., 1997; Hollerman and Schultz, 1998; Waelti et al., 2001; Bayer and Glimcher, 2005) and in the ACCs (Amiez et al., 2005; Matsumoto et al., 2007; Sallet et al., 2007). In humans, neuroimaging studies have found prediction error signals in several regions including the VTA (D'Ardenne et al., 2008) and the ACCs (Holroyd et al., 2004; Nee et al., 2011). In addition, the striatum and the OFC (O'Doherty et al., 2001; McClure et al., 2003; O'Doherty et al., 2003; Ramnani et al., 2004b; Seymour et al., 2004; Amiez et al., 2005; Bray and O'Doherty, 2007; Matsumoto et al., 2007; Sallet et al., 2007; Seymour et al., 2007; Brovelli et al., 2008; Hare et al., 2008), which are both connected to the ACCs and the VTA, show prediction error responses for different forms of prediction and outcome. It has been suggested that the nature of the prediction error signal is dependent on task context (O'Doherty et al., 2004; Seymour et al., 2007; Burke et al., 2010; Glascher et al., 2010). Thus, the brain area that codes a prediction error signal may depend on the nature of the prediction and the form of learning that is taking place. Interestingly, there is some evidence that in the ACCs these signals code for the discrepancy between the predicted value of a chosen actions and its actual outcome (Matsumoto et al., 2007). I propose that prediction error signals in the ACCg code for the discrepancy between the value that another is placing on an action and its

actual value known by the subject. When teaching, such a discrepancy may occur when a teacher observes a student's response. This chapter investigates whether prediction error signals occur in the ACCg at the moment in time that a teacher observes a student's response.

Fifteen subjects took part in an fMRI experiment conducted in two phases. In the first-phase, a training session, the subjects learnt the arbitrary associations between 10 instruction cues and one of four possible responses on a keypad. On the following day, subjects took on the role of a teacher whilst inside the MRI scanner. They observed the responses of another naïve participant (a confederate) who was tasked with learning the same associations that the teacher (participant) had learnt during training. The subject was required to monitor the responses of this 'student' and provide them with error feedback. A Rescorla-Wagner (R-W) (Rescorla and Wagner, 1972) learning algorithm was fitted to the behaviour of the student. Activity time-locked to the cues which signalled the responses of the student to the teacher were examined. The parameters from the learning algorithm were fitted to this point in time. This allowed for testing of the hypothesis that the ACCg would signal the erroneous predictions of the student, in the brain of the teacher, in the manner predicted by RLT.

4.3 Methods

Subjects were sixteen healthy right-handed participants (aged between 18 and 30; 10 female), screened for psychological and neurological disorders. One subject failed to complete the whole scanning session and was excluded from the analyses. Subjects were paired up with one of three confederate participants, who they believed were a naïve participant. The subjects were not paid for their participation but were given the incentive of receiving a picture of their brain for taking part. The subjects were informed that the other participant performing the task with them (confederates) was being paid £5 for their participation as they were not being scanned.

4.3.1 Experimental design

4.3.1.1 Training

Subjects were trained in two phases one day prior to scanning. In the first phase, the subject was seated in front of a monitor, with a response keypad. They were presented with a series of visual stimuli on the screen. Each trial consisted of an instruction cue (a coloured shape), a trigger cue (four white lines which indicated that a response should be made), the response cue (an asterisk over one of the lines indicating which response had been made) and feedback (a one pound sterling coin indicating a correct response or a one pound coin with a cross through it, indicating an incorrect response). They were required to learn, through trial and error, the arbitrary stimulus-response associations between ten instruction cues (coloured shapes that gave no indication of which response was correct) and one of four motor responses. There were 100 trials in total, with ten presentations of each instruction cue. The instruction cues were presented in two blocks, five instruction cues in the first 50 trials and five in the last 50 trials (the rationale for this is provided in 4.3.1.2). The cues were pseudorandomly presented in each of these blocks. If the subjects did not respond within 750ms of the trigger cue, feedback was displayed as “missed”. This first phase of the training ensured that all subjects had learnt all the stimulus-response associations. All subjects performed at least two contiguous, correct responses for each instruction cue. This enabled them to act as a teacher during the scanning session.

In the second phase of the training session, subjects practised the task that they would perform in the scanner on the following day (see task design below). However, rather than performing the task with a confederate as they would during scanning, they practised the task with the experimenter. This enabled the subject to become proficient at providing feedback. It was not possible for them to provide feedback to the confederate student during training. If the student learnt associations during training, they would no longer be naïve to the associations during the scanning session. It was crucial for the design that the subject believed that the student had not learnt the correct associations before the scanning session. Both the experimenter's responses during this phase of training and the confederate's responses in the scanning session, were actually a set of pre-programmed computer-controlled responses.

4.3.1.2 Scanning Session (see fig.4.1)

During the scanning session the subject acted as a teacher and taught the associations they had learnt during training to the confederate student. They observed the trials and responses of the student in real-time. The subject's task was to monitor the responses of the student and provide error feedback. Before the subject entered the scanner they were shown the student sitting in front of a monitor the corner of which was clearly covered. This enabled the teacher to be presented with stimuli that were not presented to the student. On each trial the subjects were presented with a cue that indicated the correct response ('the association cue'; a number from one to four, corresponding to the buttons on the keypad). This information was presented in the corner of the screen and was therefore not available to the Student. This cue was presented to remind the teachers of the correct associations they had learnt one day previously. The teachers indicated their feedback by making one of two responses on a keypad at the time of a trigger cue. This teacher trigger cue was also presented in the corner of the screen, ensuring that the student was not aware of the feedback they would receive before the feedback cue at the end of trial. The trials were in the same order of presentation as those during the first training phase. Thus, there were 100 trials, five associations were learnt in the first 50 trials and five in the second 50 trials. Learning was separated into two blocks to ensure that learning related activity occurred throughout the session and not just in the trials at the beginning of the session.

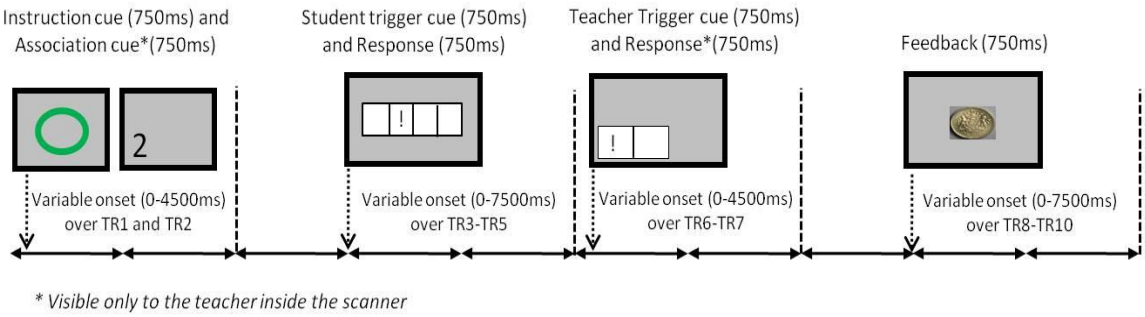

Fig.4.1 Trial Structure. Each trial started with an instruction cue (a green shape, the form of which provided no information as to the correct response); an association cue which informed the teacher of the correct response for the instruction cue; the student's response, a trigger cue instructing the teacher to make a response and a feedback cue. The association cue and the teacher's trigger cue and response, were displayed in the corner of the screen and were therefore not visible to the student.

4.3.2 Trial structure

The stimuli used during the scanning session were the same as those the teacher had observed during the training session. The teacher believed that the student, seated in the control room next to the scanner, was learning via the feedback in a similar manner as they had during the training session. However, in fact, these trials were computer controlled and the responses made were actually those of a participant during a pilot experiment. This participant's behaviour was selected on the basis of their learning rate being close to the average learning rate of the 6 participants during the pilot experiment. It was noted during this pilot experiment that there is sometimes a tendency for subjects to learn by selecting buttons sequentially from left to right in the keypad for each stimulus. This could potentially confound the results of this study, as the design relied on the subject not being able to predict the actions of the student. If the teacher could predict the action of the student, then an additional prediction error, other than that which is hypothesised, might occur at the time of the student's response. That is, the teacher may predict the action that will be taken by the student, and update this prediction through a prediction error signal when the student response occurs. Thus, the computer-controlled behaviour of the student in this experiment did not exhibit any systematic or strategic response patterns.

The feedback in this session was provided by the teacher and not by the computer. The teacher's trials consisted of an instruction cue (one of the ten they had learnt associations for

during training), immediately followed by an association cue (informing the subject of the correct association for that instruction cue), a student trigger cue (instructing the student to make a response), the student response (indicating which response the student had made), a teacher trigger cue (instructing the teacher to make a response), the teacher response cue (displaying the response of the teacher) and the feedback (indicating to the student whether the response was correct or incorrect). The subject believed that the student was responding to these trials in real-time. This subject also did not make strategic responses in order to learn these associations.

4.3.3 Computational Modelling

4.3.3.1 Behavioural Modelling

In this study, the learning of the student was modelled using a R-W algorithm (Rescorla and Wagner, 1972), using the same approach as that of Brovelli et al., (2008) who used it to model first-person learning of arbitrary visuomotor associations (Brovelli et al., 2008). Recently, more sophisticated algorithms have been used to model reinforcement learning processes, proving more accurate in predicting decision-making behaviours (Behrens et al., 2007; Behrens et al., 2008; Dayan and Daw, 2008). However, there is still some debate as to whether such approaches are better for modelling all types of learning (Dayan and Daw, 2008). In contrast, there is a considerable body of evidence which shows that the R-W model provides an accurate account of the learning of arbitrary visuomotor associations (Schultz et al., 1997; Schultz, 2006). More recently, it has also been used to explain how visuomotor associations are learnt though observation during social interactions (Burke et al., 2010). Thus, given the evidence to suggest that the R-W model provides an accurate account of the type of learning that was being performed by the student in this study, this model was employed here.

Another important consideration for studies using computational approaches is the importance of comparing plausible alternative models which could explain both behaviour and brain activity. It has been suggested that without comparing alternative models, one cannot make inferences about the precise nature of the computations that are being performed in the brain (Mars et al., in press). Thus, it has also been argued that model comparison is essential for studies which aim to investigate the computations which are being performed in the brain. However, I argue that in certain experimental paradigms, model comparison may not be

informative for the hypothesis under investigation and also, as in the case of this study, may not provide any further information as to the computations that are performed in the brain. I make this case for two reasons. Firstly, Mars et al. (in press) argue that model comparison is necessary based on assumptions about the aims of computational studies. Their assumption is that model-based studies are aiming to (i) identify whether any brain area codes a particular value or parameter from a computational model and (ii) to best explain the algorithms computed in the brain which drive the observed behaviour. However, in this chapter, the aims of the study contrast with those assumed by Mars et al. (in press) The aims were to (i) examine whether a particular computation was performed in a specific brain area and (ii) to examine whether activity in this area could be explained by a model that conformed to the basic principles of a theory of learning. Thus one of the aims of this study was to identify whether activity in the ACCg conformed to the basic principles of RLT and not to accurately define the computation being performed in this area. Secondly, in this study, the model is not fitted to the behaviour of the subjects themselves. The model is fitted to the behaviour of the student, which is in fact a set of pre-programmed responses based on the behaviour of one subject. Thus, the behaviour observed by every scanned subject (teacher) and the behaviour that the model is fitted to is identical. It therefore follows that one computational model will always fit the behaviour of the student best. Thus, model comparison at the behavioural level would only be informative for inferring the algorithms that drove the behaviour of the student. This would not be informative for inferring how the brain of a teacher models the learning of a student. This study therefore used a well established model that has been used to model conditional motor learning behaviour previously, the Rescorla-Wagner algorithm (Rescorla and Wagner, 1972; Brovelli et al., 2008). This model conforms to the computational principles of RLT and therefore is useful for testing the hypotheses that were outline in section 4.2.

4.3.3.2 The Rescorla-Wagner algorithm

The R-W model assumes that the associative value of an action (or stimulus) changes once new information reveals that the actual outcome of a decision is different from the predicted outcome (Rescorla and Wagner, 1972). Thus, on each trial, an action has a predicted associative value that is updated by a prediction error signal when an outcome reveals that this prediction is erroneous. The evolution of the associative values for each action are given by:

(1)

Where:

(2)

In both (1) and (2), *n* is the trial number, *a* = 1 ….*k* with *k* being the number of available actions and η is the learning rate. The asymptotic value (λ) of a correct action increases to greater than 0 once a correct response has been made and is 0 for an incorrect response. The prediction error is therefore the student's prediction of its associative value (subtracted from the actual value of the action ( ) known by the teacher. We instructed the students (and teachers on the first day) that 1 of the four finger movements could be correct for each stimulus. Importantly, this also ensured that learning the correct association for one instruction cue was not informative as to the correct associations for any other instruction cue. Thus the associative values of actions for one instruction cue were not informative as to the value of an action for another instruction cue. The initial associative strength of each action for each stimulus was set to 0.25, given the equiprobability of each of the four actions being correct.

4.3.3.3 Model estimation

To model the action selection process of the student the associative values were transformed into probabilities using the softmax equation (see (2) below). These probabilities reflected the likelihood of the actions actually chosen by the student, given the value of the free parameters in the model. Thus, by varying the value of the free parameters, one can find the value of the parameters that maximise the probabilities of the actions actually chosen. The values of the parameters that result in the highest probabilities in this algorithm are assumed to reflect the

processes underpinning the behaviour of the subject during learning. This method is a standard approach used in RLT (Sutton and Barto, 1981). The probability of the action chosen by a subject is given by:

(3)

This equation converts the associative values of the action chosen by a subject to a probability (. The coefficient β represents the stochosticity (or temperature) of the student's behaviour (i.e. the sensitivity to the value of each option). A high β (greater than 1) results in the performance of all actions becoming nearly equiprobable, with a low β amplifying the differences in associative values. These two algorithms were used to model the action selection of the student over-time. The associative value that the student placed on the chosen action ( ) was then updated in the R-W model, based on the feedback. Crucially, in this study, this feedback was provided by a teacher inside the scanner. As the teacher had expert knowledge of all the associations –and was informed of the correct action on each trial- they knew the asymptotic value (λ) of each action chosen by the student. In this experiment, the aim was to examine whether the teacher modelled the learning of the student. It was therefore assumed that to instruct the student, the teacher would have to calculate the discrepancy between the student's prediction of the outcome and the asymptotic value (λ) of the action chosen by the student. This asymptotic value is known only by the teacher whilst the student is still learning; only once the student has learnt the correct stimulus-response associations for each cue will there be no discrepancy between the asymptotic value known by the teacher and the prediction made by the student. The aim of the teacher is therefore to provide the student with appropriate feedback to minimise the discrepancy between their expert knowledge and predictions made by the student.

Within the R-W model and the softmax algorithm there are free parameters which need to be estimated. To identify the optimal set of free parameters for the student's behaviour (given the teacher's feedback), the learning rate, the stochasticity parameter β and the asymptotic value λ were varied. The output of the softmax algorithm is a series of probability values, based on the value of each of the free parameters and the actions chosen by the student. When the value of the free parameters are varied, the probability values output by the softmax algorithm change. To select the parameters that best fitted the student's behavioural

data (given the teacher's feedback) a maximum likelihood approach was used. By using a maximum likelihood algorithm it was possible to maximise the probabilities of the actions chosen by the student and to identify the values of the free parameters which produced the highest probabilities. The learning rate (η) was varied between 0 and 1 in steps of 0.05, β between 0 and 5 in steps of 0.1 and λ between 0 and 5 in steps of 0.1. The likelihood of the chosen actions were found using:

(4) (n)

Where the likelihood of each set of parameters (L) is determined by the log of probability of the performed action ((n)) of the student at trial *n,* according to the model. If the model perfectly predicts the actions, the probability of every chosen action would = 1 and L would be 0. As the probabilities become less than 1 the log-likelihood L assumes negative values. The best fitting parameters were then selected using:

(5)

This identified the set of parameters for which L was closest to 0 i.e. the best fitting parameter set. Where is the parameter set and L is the log-likelihood. Importantly, in this study, the student's data was computer controlled and thus every teacher observed the same responses of the student. Variations in these parameters could therefore only be explained by changes in the feedback, i.e. if the teacher failed to give the student feedback on a particular trial. The maximum likelihood approach revealed, that for the behaviour of the student, the best fitting parameters were a λ of 1 a learning rate η of 0.95 and a β values ranging from 2.3 to 2.7- reflecting the differences in stochasticity of the behaviour given the teacher's feedback.

4.3.4 Functional imaging and analysis

4.3.4.1 Data acquisition

1016 EPI scans were acquired from each participant. 38 slices were acquired in an ascending manner, at an oblique angle (≈30˚) to the AC-PC line to decrease the impact of susceptibility artefact in the subgenual ACC (Deichmann et al., 2003). A voxel size of 3×3×3 mm (20% slice gap, 0.6 mm) was used; TR=3 s, TE=32, flip angle=85°. The functional sequence lasted 51 minutes. High resolution T1-weighted structural images were also acquired at a resolution of 1×1×1 mm using an MPRAGE sequence. Immediately following the functional sequence, phase and magnitude maps were collected using a GRE field map sequence (TE_1 = 5.19ms, TE_2 = 7.65ms).

4.3.4.2 Image preprocessing

Scans were pre-processed using SPM8 (www.fil.ion.ucl.ac.uk/spm). The EPI images from each subject were corrected for distortions caused by susceptibility-induced field inhomogeneities using the FieldMap toolbox (Andersson et al., 2001). This approach corrects for both static distortions and changes in these distortions attributable to head motion (Hutton et al., 2002). The static distortions were calculated using the *B*0 field map acquired after the EPI sequence. The EPI images were then realigned, and coregistered to the subject's own anatomical image. The structural image was processed using a unified segmentation procedure combining segmentation, bias correction, and spatial normalization to the MNI template (Ashburner and Friston, 2005); the same normalization parameters were then used to normalize the EPI images. Lastly, a Gaussian kernel of 8 mm FWHM was applied to spatially smooth the images in order to conform to the assumptions of the GLM implemented in SPM8.

4.3.4.3 Event definition and modelling (Student response)

In this study, multiple GLM analyses were performed to investigate activity time-locked to the teacher's observation of the student's response. These were performed to ensure that activations identified could only be accounted for by the independently explained variance of a parameter in the R-W model. Although each of the GLMs differed from the others, they shared

several common properties. Each GLM analysis contained regressors modelling the Instruction cue event, the student response cue, the teacher trigger cue and the feedback cue. Regressors were constructed for each of these events by convolving the event timings with the canonical Heamodynamic Response Function (HRF). The residual effects of head motion were modelled in the analysis by including the six parameters of head motion acquired during preprocessing as covariates of no interest. In addition to these regressors defined for the event types, each GLM also contained regressors which were first order parametric modulations of the student response cue event. These modulators scaled the amplitude of the HRF in line with either the λ_a , V_a or δ parameters from the R-W algorithm. The values of these parameters corresponded respectively to the Teacher's valuation; the Student's prediction and the student's prediction error, known only by the teacher at that time (see fig.4.2 for example model parameters). When a trial was missed by the student, these parameters all took on a value of zero. Two sets of analyses were conducted in this study to examine responses in the brain of the teacher at the time of the student's response:

(1) Three separate GLMs were created in which the values of one of λ, V_a, and δ were used as first-order parametric modulators of the student response cues. These models enabled areas of the brain in which the BOLD response varied in the manner predicted by one of the parameters to be identified. However, due to correlations between the values of these parameters in the R-W model, additional analyses were conducted. These GLMs orthogonalized one of the parameters with respect to another. Thus, voxels in which activity varied in the manner predicted by the orthogonalized regressor, explained a portion of the variance uniquely. For these analyses, two GLMs were created for each of the three parameters. Within these GLMs the regressor of interest was orthogonalized with respect to one regressor modelling one of the other parameters of the model. For example, for the δ parameter, three GLMs were constructed in total. The first contained only the values of δ parameter as a parametric modulation of the student response cues. The second contained λ, as a parametric modulator, with the values of the δ parametric modulator orthogonalized with respect to the values λ. The third contained Va, as a parametric modulator, with the values of the δ parametric modulator orthogonalized with respect to the values of the Va parameter. Regressors were orthogonalized in SPM using a Gram-Schmidt process. Parametric modulators are orthogonalized serially by their position in the design matrix, such that the second parametric regressor is orthogonalized with respect to the first. A *t*-contrast was performed on each of these regressors in the separate GLMs and taken to a second-level. Voxels are only reported if they were significant in an *F*-contrast in all three of these GLMs at the second-level.

This approach was then repeated for the λ and V_a parameters. Thus, nine GLMs were constructed to examine activity which varied with the values from the parameters of the R-W model. It is important to note that typically one would orthogonalize the parameter of interest with respect to both of the other parameters, in one GLM. However, this was not possible in this study, as the δ parameter was a product of the other two parameters in the R-W model. Thus, orthogonalizing the δ parameter with respect to both of the other parameters in this model would have removed all statistical power from the regressor.

(2) An additional GLM was created to control for other possible responses in the ACC at the time of the student's response. This GLM contained control parameters that varied with other plausible responses which were not a component of the R-W model. Firstly, it is possible that the ACCg could compute the magnitude of the prediction error response, rather than the signed prediction error response as in the R-W model. Thus, an 'unsigned' parameter was constructed that contained the values of the magnitude of prediction error response, i.e. it contained only positive values. In addition, it was also possible that the ACCg could signal the prediction error responses in a non-parametric manner. To test whether the ACCg made a non-parametric response whenever the prediction of another was erroneous, a parameter was included that took on a value of 1 when there was an error and 0 when there was no error. These parameters were fitted to the responses of the student and included in a GLM, in which they were not orthogonalized with respect to each other. *t*-tests were then conducted between them to test which model best explained activity in a given voxel.

4.3.4.4 Event definition and modelling (Feedback)

In addition to the analyses of activity at the time of the student response, analyses were also conducted on the feedback cue activity. At this point in each trial, the student is able to calculate the discrepancy between their prediction and the outcome. Thus, if any of the brain codes a representation of another's prediction error signal (i.e. an area that responds at the exact time that they see another experiencing a prediction error) then activity in such an area would vary with the δ parameter in the brain of a teacher at the time that the feedback is given to the student. The strategy employed for this analysis followed those employed in the first set of analyses. However, for these analyses all parameters were fitted to the feedback cues rather than the student's response cues.

4.3.4.5 Group analysis, Contrasts and Thresholding

Random effects analyses (Full-Factorial ANOVA) were applied to determine voxels significantly different at the group level. SPM{t} images from all subjects at the first-level were input into second-level full factorial design matrices. *t*-contrasts and *F*-contrasts were conducted in each of the GLMs. These contrasts identified voxels in which activity varied parametrically in the manner predicted by the parameters in the R-W model. Separate corrections for multiple comparisons were used for the ACCg and the whole brain. To examine activity across the whole brain, FDR correction was applied. In contrast, activity in the ACCg was corrected for by using an 80% probability mask of the ACCg (see general methods for a description).

For the second set of analyses examining activity time-locked to the student's response cues, the *t*-contrasts between the prediction error parameter and the control parameters were examined at a lower threshold. This was necessary due to the high correlations between each of these parameters. For these contrasts a threshold of $P<0.01$, uncorrected for multiple comparisons, was employed.

4.4 Results

4.4.1 Behavioural results

The teacher's task was to monitor the student's responses and determine whether each response was correct or incorrect. The student's responses, unbeknown to the teachers, were a computer controlled replay of a real subject's responses during a pilot experiment. The student missed three trials and thus, teachers were required to respond on 97 trials. Subjects correctly gave feedback to the student on 95.2% (SD ± 2.9; range: 91-99%) of trials, indicating that all subjects understood the correct association for each stimulus and also understood whether the student's responses were correct or incorrect.

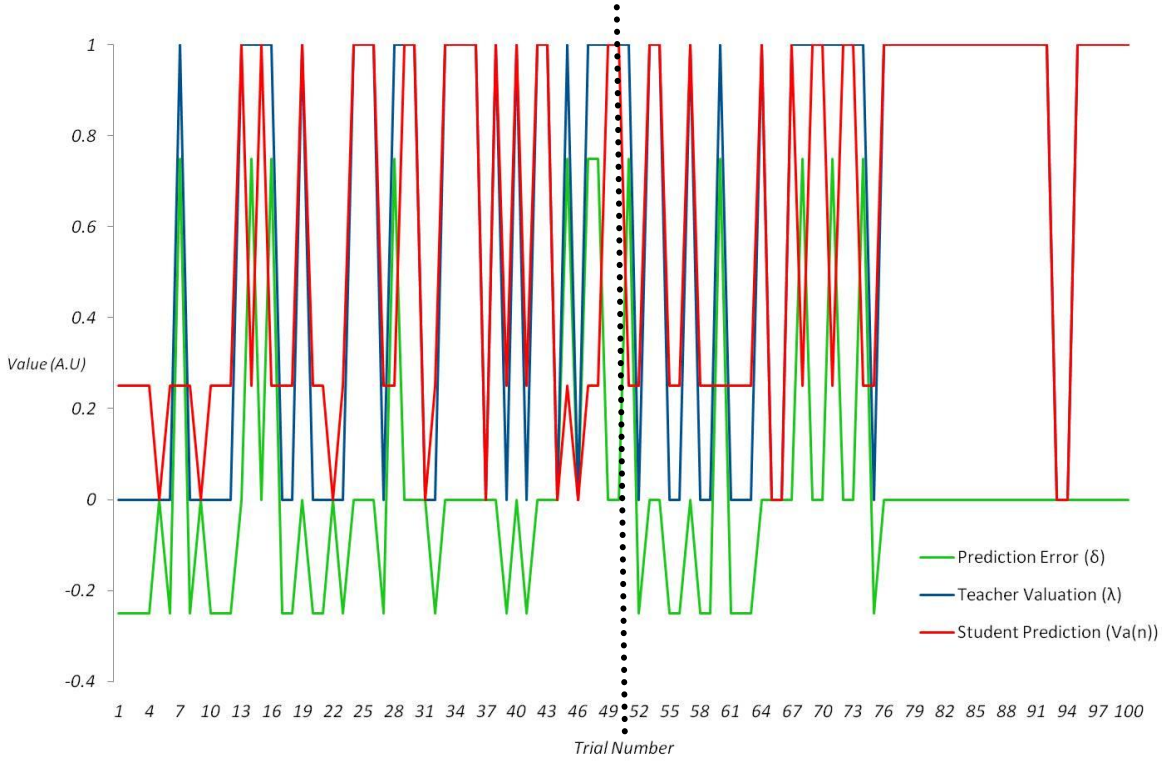

Fig. 4.2 Example parameters from the R-W model. Each of these parameters were then fitted to the moment in time at which the teacher saw the student's response. Dotted line indicates the point in which a new set of instruction cues was introduced.

4.4.2 Imaging results

4.4.2.1 Prediction error(δ)

The main aim of this experiment was to examine prediction error responses at the moment in time when the teacher saw the student's response. This study tested the hypothesis that the ACCg would signal the discrepancy between a student's prediction and the actual outcome known by a student. In line with the hypothesis, activity was found in the ACCg which varied significantly with the prediction error (δ) parameter of the R-W model (4, 30, 12; F(1,84) = 8.49 Z = 3.17, p<0.05svc). Activity in this area was fitted marginally significantly better by the prediction error parameter from the model, than either the unsigned or non-parametric control regressors (p<0.05 uncorrected (unc)). In addition, no area of the brain showed a significantly greater response to either of these regressors than the prediction error parameter from the R-W model. No other brain area significantly varied with the prediction error parameter. Two other clusters in the VTA and the Caudate did not survive correction for multiple comparisons, but were the only regions activated even at a much lower threshold (p<0.05 unc).

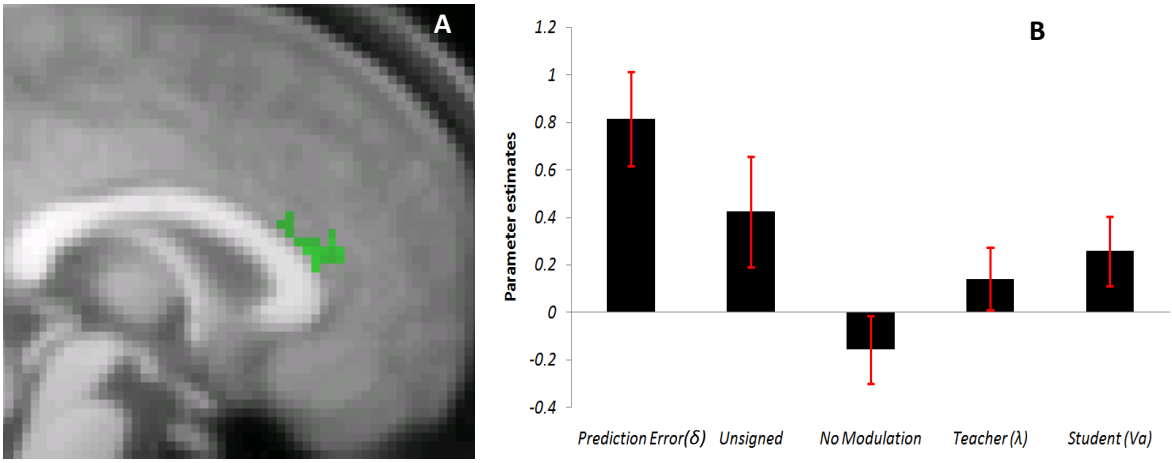

Fig 4.3.Prediction error signal in the ACC. (A) Activity displayed on a mean, normalised anatomical image. Activity in the ACC covaried with the prediction error parameter estimated in the model. (B) Parameter estimates (Beta coefficients) for regressors of in the peak ACC voxel. The graph clearly shows that the prediction error parameter fitted the BOLD response in this area better than the actual value of the action ( ) or the student's prediction . The beta coefficients of two additional regressors are included; one regressor modelled the unsigned magnitude of the discrepancy between the teacher's valuation and the student's prediction ('unsigned'). The second modelled a non-parametric response to every incorrect and first correct trial ('no modulation'). The prediction error parameter fit the BOLD response significantly better than the unsigned parameter (F(1,112) = 9.24; Z = 2.75;p<0.01unc) and also the unmodulated parameter (F(1,112) = 12.38; Z =3.23; p<0.05FDR). The beta coefficients of the prediction errpr parameter are displayed once the regressor had been orthorgonalized with respect to student's prediction parameter.

4.4.2.2 Student Prediction

In addition to the prediction error parameter, the student's prediction ( ) parameter was also fitted to the moments in time when the teacher saw the student's response. Activity which varied statistically with this parameter was found in a ventromedial prefrontal area (VmPFC) extending over gyrus rectus, the superorbital sulcus, and the medial superior frontal gyrus (across areas 32, 10p and 10r) and bilaterally in the insula (putatively area Idg).

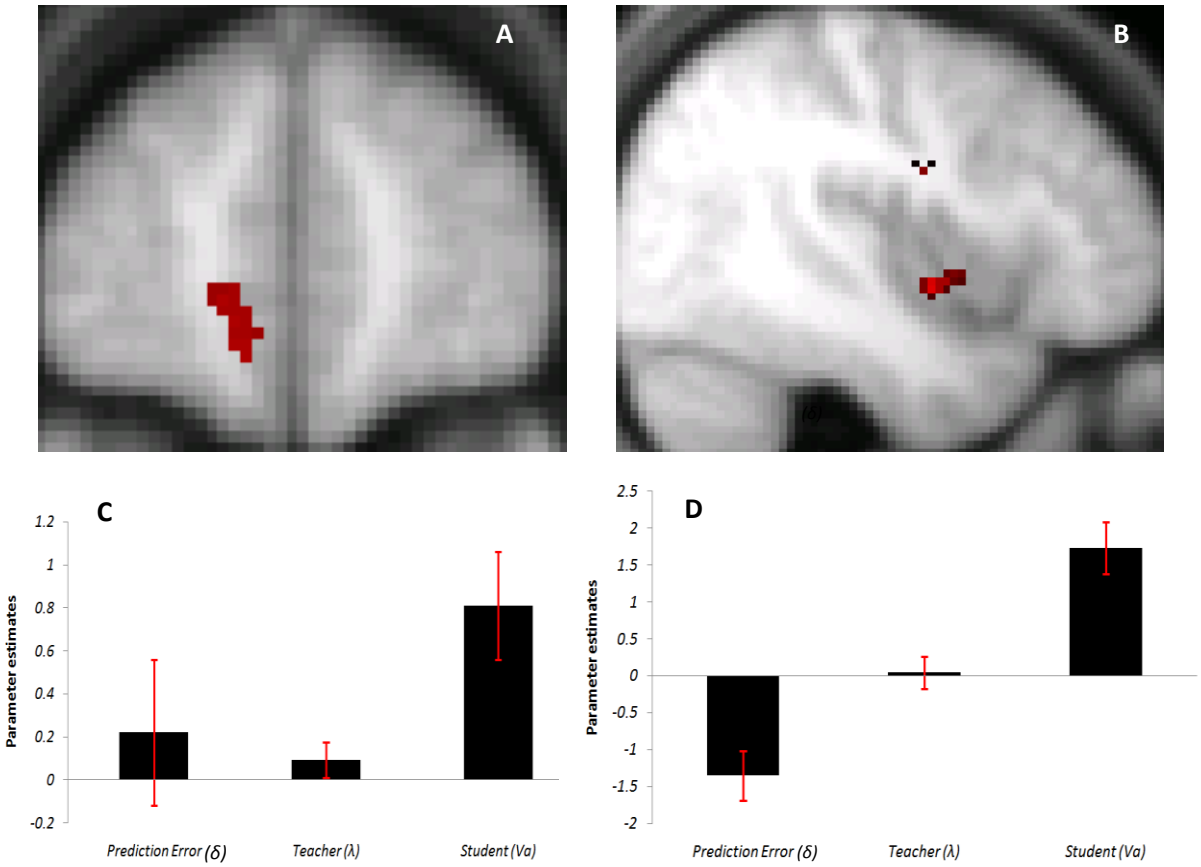

Fig 4.4 Student's Prediction. BOLD responses which varied with the values of the student's prediction parameter were found in two regions; the Ventromedial Prefrontal Cortex (A; -14,32,-10; t(84) = 4.61 Z= 4.34, p<0.05 FDR corrected for multiple comparisons) and bilaterally in the Insula cortex. Activity in the right insula (48,-4,-2;t(84) = 4.14; Z=3.94, p<0.05 FDR corrected) is reported in B and D. Plots of the parameter estimates for each of the parameters of the model are show for VmPFC (C) and insula (D). The beta coefficients of the Student prediction parameter are displayed once the regressor had been orthorgonalized with respect to the teacher's valuation parameter (the regressor with which it shared the most variance). The other Betas reflect the parameter estimates without being made orthogonal to any other parameter.

4.4.2.3 Teacher Valuation

Finally the asymptotic value of the action (λ), known only by the teacher, was fitted to the student's responses. Activity which varied statistically with this parameter (see fig. 5.4 was found in the Superior Frontal Sulcus (SFS) bordering areas 8, 9 and 9/46 and Posterior Cingulate Cortex (PCC; putatively BA 23).

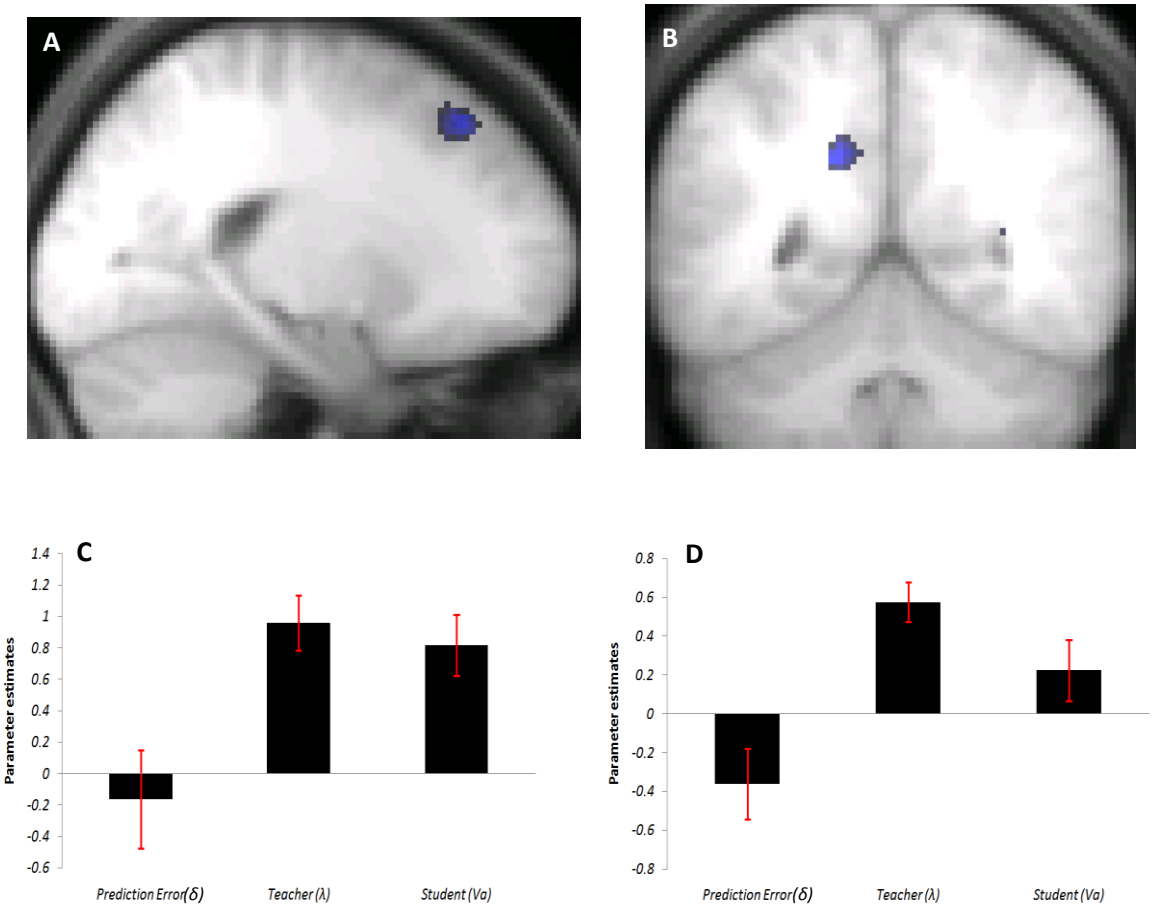

Fig.4.5 Teacher's valuation. BOLD responses which varied with the true value of the action known only by the teacher. Activity varied with the values of this parameter in only two regions, the Posterior Cingulate Cortex (A; -14,-52,32; t(84) = 5.84; Z=5.34, p<0.05 FDR) and the Superior Frontal Sulcus (B; -20,32,46; t(84) = 5.07; Z = 4.72;p<0.05 FDR). Plots of the parameter estimates for each of the parameters of the model are shown for the PCC (C) and the SFS (D). The beta coefficients of the teacher's valuation parameter are displayed once the regressor had been orthorgonalised with respect to the student's prediction parameter (the regressor with which it shared the most variance). The other Betas reflect the parameter estimates without being made orthogonal to any other

4.4.2.4 Feedback-related Activity

The same analysis strategy was employed for examining activity at the time of the feedback cues as that which was employed for examining activity at the time of the student' response. Activity was not found to vary significantly with any of the parameters in any brain area.

4.5 Discussion

This study investigated activity in the brain of a teacher when monitoring the responses of a student. It was predicted that activity in the ACCg would vary with the discrepancy between the student's prediction and the actual outcome of a trial. In line with this hypothesis, activity in a portion of the ACCg, in midcingulate area 24b', varied in the manner predicted by the prediction error parameter of the R-W model. In addition, activity was found in other areas which covaried with the other parameters in the R-W model. Activity in the SFS and the PCC varied with the Teacher's valuation of the outcome and activity bilaterally in insula cortex and in the VmPFC varied with the student's prediction of the outcome. These results support the notion that teacher's model the learning of students in a manner that conforms to the principles of Reinforcement Learning Theory (RLT).

4.5.1 Prediction errors

The results of this study suggest that the ACCg calculates the discrepancy between another's prediction and the actual outcome of their decision. Interestingly, each of the areas in which activity varied with one of the parameters of the R-W model, show reciprocal connections to the ACCg (Pandya et al., 1981; Vogt and Pandya, 1987; Carmichael and Price, 1995; Petrides and Pandya, 2006, 2007). Given that the ACCg has access to information processed in areas of the brain which are coding the predictions and outcomes of others' decisions, it is well-placed to code for the discrepancy between their prediction and outcome.

This study is not the first to investigate the processing of errors in the ACC in a social context. In the ACC an Error-Related Negativity (ERN) signal occurs when subject's experience an error (Frank et al., 2005; Holroyd et al., 2009). When these errors pertain to the actions, decisions and outcomes of others' behaviour, the signal is attenuated (van Schie et al., 2004; Bellebaum et al., 2010; Kang et al., 2010; Koban et al., 2010). Previously, it has been suggested that this attenuation of the signal was a result of changes in the nature of the signal evoked in the ACC, i.e. the same portion of the ACC makes a different response when monitoring other's errors. However, the results reported in this study and in the previous chapter suggest that the attenuation may be the result of error signals occurring in the ACCg when processing another's prediction errors. The limited spatial resolution of EEG would make it difficult to disentangle sources originating in the sulcus from those originating in the gyrus. Thus, the attenuation of

the ERN may be a result of additional activations occurring in the ACCg when errors pertain to the behaviour of others.

Neuroimaging research has also investigated the processing of errors in a social context and, as stated in the introduction, some of these have been in the context of social learning. In these studies, the erroneous actions of other's have been shown to activate the ACC (Shane et al., 2008; Shane et al., 2009), the motor system (Malfait et al., 2010), the striatum (King-Casas et al., 2005; Burke et al., 2010) and also the pSTS and paracingulate cortices (Hampton et al., 2008). In this study, a prediction error signal was found in the ACCg. This result and those listed above, support the view outlined in chapter one, that different areas of the brain process different forms of prediction error signal. In this study, the results indicate that the ACCg codes for the discrepancy between the prediction of another and the outcome that they will receive.

The design of this study had several advantages over those investigating error processing in a social context. Firstly, in this study, activity was time-locked to the exact moment that the response of the student was observed by the teacher. Previous studies investigating the processing of other's errors have examined activity at the time of another's response (Shane et al., 2008; de Bruijn et al., 2009; Shane et al., 2009; Kang et al., 2010). However, these studies have not examined how the erroneous actions of another are processed when the subject is engaged in monitoring and providing another with feedback. Secondly, this study used a R-W learning algorithm to model the learning behaviour of the student and test whether activity in the brain of a teacher covaried with the parameters of this model. Previous studies have used similar approaches to examine whether the learning of another is modelled in the brain (Behrens et al., 2008; Hampton et al., 2008; Burke et al., 2010). However, in these studies, the subjects own learning was contingent on the behaviour of the other individual. Thus, the prediction errors pertained to the subjects own predictions about the other agent being erroneous. These studies therefore did not examine whether prediction error signals occur when it is another's prediction that is erroneous and not the subjects own prediction. This study therefore supports the interpretation of prediction error signals in the ACCg that was outlined in the previous chapter. That is, the ACCg signals the erroneous predictions of another in a manner that conforms to RLT. It therefore highlights how prediction error signals occur in a distributed network of brain areas, in which the network that processes the prediction error signal is dependent on the context of the task.

4.5.2 Processing others' predictions

It is not surprising that activity was found in superior portion of the OFC, in an area often referred to as Ventromedial Prefrontal Cortex (VmPFC), and in the insula varied with the student's prediction parameter given that these two areas exhibit strong reciprocal connections (Mesulam and Mufson, 1982; Mufson and Mesulam, 1982). In addition to connections to the insula, the VmPFC shows strong connections to the ventral striatum (Yeterian and Pandya, 1991), ACCs, and medial and lateral parts of the OFC (Pandya et al., 1981). These areas are all well-known for their roles in processing rewarding stimuli. This connectional fingerprint would therefore implicate a portion of VmPFC in processing reward-related information. Neuroimaging investigations in humans support this notion, with activity in the VmPFC being modulated by expectations about rewarding outcomes. More specifically, the BOLD response in this area is scaled by the magnitude of an expected reward following choice selection (Hampton et al., 2006; Chib et al., 2009; Hare et al., 2009; Wunderlich et al., 2009; Prevost et al., 2010; Tricomi et al., 2010). Thus, it appears that the VmPFC codes the predicted reward following a choice. Behrens et al., (2008) also found that the activity in this area reflected the reward value that a subject expected to receive. However, they showed that in some subjects, activity in this area coded the expected reward value, based on the advice given by another. This would suggest that the expected value signal in the VmPFC is also sensitive to the expected outcomes of other's decisions. In this study, like in Behrens et al (2008), activity in the VmPFC was scaled by the predicted reward value of another. However, in this study, the subjects were not expecting a rewarding outcome themselves. As such, the VmPFC is coding the subjects' representation of the student's expectations. This would suggest that the VmPFC can flexibly code expected reward values, when they pertain either to one's own or another's predictions.

Activity in the insula was also scaled by the student's prediction of the outcome. Area Idg (putatively the area activated in this study), in the insula has widespread connections across the cortex (Mesulam and Mufson, 1982; Mufson and Mesulam, 1982). As a result, it has been implicated in a large number of processes including the processing of rewards (Knutson et al., 2000; Knutson and Bossaerts, 2007; Preuschoff et al., 2008), risk (Cardinal, 2006; Preuschoff et al., 2008) pain (Singer et al., 2004) and emotional processes such as fear (Adolphs, 2002; Phan et al., 2002). There is also evidence that this area may be involved in the processing of social information (Singer et al., 2004; Jabbi et al., 2008; Rilling et al., 2008; Baumgartner et al., 2009; Singer et al., 2009). However, the number of processes which activate this portion of the insula

and the extent of its connections across the cortex, has resulted in absence of a coherent picture of its functional properties. Despite the absence of a precise theoretical account of information processing performed in Idg, recent research has suggested that information is processed in this area in a manner that conforms to the principles of RLT (Singer et al., 2009). That is, within the insula, activity is scaled by different forms of predictions, such as the predicted amount of pain, or predicted level of risk. When new information reveals that the predicted values were erroneous, error signals occur, updating future estimates of the predictions. Such prediction and prediction error signals have been shown for different types of information in the insula, including risk (Preuschoff et al., 2008; Quartz, 2009) and uncertainty (Rolls et al., 2008). Thus, the insula codes information in a manner that conforms to the principles of RLT (Rushworth et al., 2009; Singer et al., 2009). However, these predictions are distinct from typical reward prediction errors. In this area the predictions and therefore error are reward-related but modulated by a particular decision-making variable such are risk or uncertainty. In this study, reward, risk and uncertainty were not manipulated, as they were not important for the specific hypotheses under investigation. As such, it is not possible to make inferences about whether the activity correlating with the student' predicted reward value, may have been modulated by one of the other variables, such as risk or uncertainty, that is known to modulate reward values in this area.

4.5.3 Processing others' outcomes

In addition to activity correlating with the students' predictions, activity was found in the PCC and the Superior Frontal Sulcus (SFS) that covaried with the teacher's knowledge of the actual outcome. Although these regions are reciprocally connected (Petrides and Pandya, 1999), they otherwise show distinct patterns of connectivity and respond to different types of stimuli. These distinctions are informative as to the nature of the coding in these two areas in this study.

There are several different accounts of PCC functions, including suggestions that it is part of the default mode network, a set of interconnected regions that become deactivated during tasks. However, the portion of the PCC activated in this study is anterior and superior to the portion that is considered as part of the default-mode network (Cauda et al., 2010). The portion of the PCC activated in this study is implicated in processing the value of choices. This region is heavily interconnected with the striatum (Yeterian and Pandya, 1991), the ACCs and

medial and lateral portions of the OFC (Pandya et al., 1981), areas which are engaged when processing rewarding stimuli (Haber and Knutson, 2010). Single-unit recording studies have shown that the PCC contains neurons which fire quantitatively with the value of rewarding stimuli (Hayden et al., 2008). In humans, fMRI studies have shown that activity in the PCC covaries with the subjective magnitude of rewards, rather than the objective reward value (Kable and Glimcher, 2007), at the time of a choice. This would suggest that the PCC is engaged in processing the rewarding value of the outcomes of decisions. A recent single-unit recording study in monkeys showed that neurons in this region increased their spike rate when monkeys chose a reward that was accompanied by a social stimulus, compared to when a reward of the same magnitude was chosen without a social stimulus (Heilbronner et al., 2011). This study provided the first indication that, not only do neurons in the PCC process reward values, but activity in these neurons is also sensitive to social information. In this study, activity in the PCC covaried with the asymptotic value of the action chosen by the student. However, in a conditional motor learning task, modelled with R-W model, reward value and associative values are identical. When an incorrect response is made, the reward value and the associative strength of the chosen action are both zero, when a correct response is made, they both acquire values higher than zero. Thus, given the sensitivity of neurons in this area to rewards and social value, it would suggest that in this study, the PCC appears to be coding the rewarding value of another's action.

In contrast to the PCC, the SFS shows strong connections to premotor cortex and the rostral cingulate zone (Petrides and Pandya, 1999). Neurons in the SFS are known to process action-related information at its most abstract level (Passingham et al., 2010). In order to understand the decision-making of another and provide them with appropriate feedback, it may be important to have access to abstract motor information. More specifically it may be important to code for the associative strengths of actions performed by others. Ramnani and Miall (2004) reported activity in the same portion of the SFS when subjects monitored the instruction cues of another. They found that this area was engaged only when an instrumental cue signalled that the action of another was predictable and not when an action could not be predicted. The predictable actions were learnt by arbitrary association and as such, when the action of another could be predicted, the associative value of that action could also be coded. The results of Ramnani and Miall (2004) and this study therefore suggest that the SFS may play an important role in processing the associative values of actions chosen by others.

4.5.4 Caveats and Limitations

The aims of this study required that activity in the brain of a teacher was examined at the moment in time they saw the responses of a student. Additional analyses were also conducted to examine whether any area of the brain showed prediction error responses at the time of the feedback. However, no brain area responded significantly to any of the parameters in the R-W model. At first this seems somewhat surprising. Previous studies have shown a number of different areas signal prediction error responses time-locked to the outcome of other's decisions. Indeed prediction error responses have been found in the dorsal striatum (King-Casas et al., 2005), the dorsolateral prefrontal cortex (Burke et al., 2010), the ACC (Apps et al., accepted) the paracingulate cortex and the pSTS (Behrens et al., 2008; Hampton et al., 2008). However, there are important distinctions between the nature of the feedback cues in this study and those which have previously found prediction error responses. Specifically, in those studies, the feedback cue informed the subjects themselves about the outcome of the trial. This information was not available to them before this point in time. In addition, the subjects own learning was dependent on monitoring or learning from the other's outcomes. In this chapter, the outcome of the trial was already known by the subject, as they had monitored the third-persons responses and provided them with feedback themselves. The subject was also not learning from the feedback cues. As such, the subject could perform the task without monitoring the feedback cues. This lack of saliency of the feedback cue may explain an absence of any significant activity time-locked to them.

An additional limitation to the design of this study was its inefficiency for examining activity time-locked to the instruction cues. At this point in time the student would be preparing an action and thus, the teacher may be engaged in modelling the action selection process of the student. However, in this study the instruction cues were not jittered independently from the association cues that indicated the correct response for the presented stimulus to the teacher. Thus, activity occurring time-locked to the instruction cues would be confounded by signals occurring at the time of the association cue. As such, it was not possible to examine activity that occurred during the student's preparation of a response. This approach was taken in order to ensure that the student's response cues could be jittered independently from all other cues i.e. maximising the statistical independence of the cue to which the main aims of the study pertained. However, in chapter 5, a study is reported which examines activity time-locked to the instruction cues. In that chapter, hypotheses will be tested about the processing of reward-related information at the time of instruction cues.

The final limitation to the design of this study relates to the use of the R-W algorithm to model behaviour. In this study, the behaviour of the student and the activity in several areas of the brain of the teacher were significantly well explained by parameters from the R-W algorithm. Whilst this model is one of the most well established models of learning (Schultz, 2006), there are alternative models which have also been used recently to explain learning behaviour during both social and non-social decision-making (O'Doherty et al., 2007; Mars et al., in press). In this study, a comparison between models was not performed (for reasons outlined in section 4.3) and as such, it is not possible to claim that the R-W algorithm is the best model to explain the data in this study. An important point to note however, is that although other models differ from the R-W algorithm, ubiquitous across each of the models is that values are updated when predictions are erroneous (Bar, 2007; Behrens et al., 2008; Hampton et al., 2008; Yoshida et al., 2008). As the aim of this study was to examine whether activity in the ACCg varied with the discrepancy between a prediction and an actual outcome, more advanced models may have further elucidated the precise form of the prediction error signal in the ACCg. However, as other models contain variants of the prediction error parameter, the interpretation of the results of this study would have remained the same even if one of these models was a better fit to the data than the R-W algorithm. Thus, the aim of future research should be to develop models that can account for the behaviour and activity in the brain of a teacher when monitoring a student's responses.

4.5.5 Summary

This study examined activity in the brain of a teacher, whilst they monitored the responses of a student. This chapter extended the findings of the previous chapter in which activity in the ACCg was found to signal when another's prediction of an outcome was erroneous. Here, I have shown that activity in the ACCg also signals the erroneous predictions of another at the time when their response is observed. Moreover, activity in this area varied with the prediction error parameter from a R-W algorithm. These results suggest that the ACCg signals the erroneous prediction of other's in a manner that conforms to RLT. In addition, activity in the VmPFC and insula at the time of the student's response signalled the predictions of the student and activity in the PCC and SFS signalled the actual outcome that the student would receive. Thus teachers appear to model the learning of their student's in a manner that conforms to RLT.

Chapter 5: ACCg - Signalling the effort-discounted value of others' actions

5.1 Abstract

In the previous chapter, activity was identified in the ACCg at the time of another's action that coded for the discrepancy between another's prediction and the actual value of their action. Such activity conforms to the principles of Reinforcement Learning Theory (RLT). However, RLT proposes that choices between different courses of action are guided by predictions about the value of actions. These predictions are processed at the time of cues that instruct actions or choices between different courses of action. To fully test the hypothesis that the ACCg processes the valuations that other's place on actions in a manner that conforms to RLT, it is important to examine whether activity in the ACCg codes for the predicted value of other's actions. The predicted value of an option can be modulated by the amount of effort that has to be expended in order to receive it. Does the ACCg signal the value of effort-discounted rewards that are to be performed by others? In this chapter, subjects performed trials in which they had to expend differing amounts of effort (2, 3, 8 or 12 cued button presses) in order to receive either a high (16p) or low (4p) reward. They also observed identical trials performed by a third-person. The amount of reward, the amount of effort and who was to respond on each trial was instructed by a series of 16 cues. Activity time-locked to these instruction cues was investigated. This study tested two hypotheses, firstly, that activity in the ACCs would vary with the net value of rewards on the first-person trials and secondly, that activity in the ACCg would vary with the net reward value on the third-person trials.

5.2 Introduction

The two previous chapters have shown that the ACCg signals how erroneous the predictions of others are, in a manner that conforms to the principles of RLT. These results support the notion that the ACCg processes the same information about other's actions that are processed in the ACCs about one's own actions. However, although the ACCs is well known for processing prediction error signals at the time of the outcome of one's own decisions, it is also known for coding for predictions about the value of actions at the time they are instructed (Amiez et al., 2005; Kennerley et al., 2009; Hayden and Platt, 2010). Does the ACCg signal the predicted value of other's actions in the same manner that the ACCs processes the predicted value of one's own actions? The study reported in this chapter explores this question, investigating activity in the ACC at the time of cues that signal the amount of reward available and the amount of effort that must be expended for its receipt.

Choice behaviour is driven by the rewards associated with different courses of action (Doya, 2008). However, the value of a choice can be discounted by the cost associated with performing the necessary actions to receive a reward (Charnov, 1976). As such, a cost-benefit analysis is performed on each option, with the actions that have a higher net reward value (benefit – cost) preferred over actions that have a lower net reward value. Thus, the greater the cost associated with performing the actions, the more the magnitude of the reward is discounted. As a result, one course of action becomes less likely, when a less rewarding but less costly option is an alternative.

Effort is one of the most important variables that can modulate reward values and influence choice behaviour (Charnov, 1976). The majority of animal foraging behaviour is aimed at maximising rewards, whilst minimising the amount of effort expended to receive them (Charnov, 1976; Stephens and Charnov, 1982; Bernstein et al., 1991). For example, the choice behaviour of birds reflects not only the rate at which they will be rewarded, but also the metabolic costs associated with the behaviour to receive the reward (Bautista et al., 1998; Bautista et al., 2001). Similarly rat, monkey and human choice behaviour is influenced by both the amount of reward and also the number of actions (or level of effort) required for its receipt (Rudebeck et al., 2006b; Walton et al., 2006; Botvinick et al., 2009; Kurniawan et al., 2010). Work investigating the neural antecedents of effort-related costs in animals has implicated the ventral striatum and interconnected portions of the ACC (in the rostral cingulate motor areas). Lesions to the dACC (areas Cg1 and Cg2) of the rat have been shown to bias animals away from

investing a greater amount of effort to receive a greater reward (Rudebeck et al., 2006b). Single-cell recording studies in both rats and non-human primates have found neurons in the ACC (putatively in ACCs area 24c in primates) in which activity is sensitive to both reward magnitude and also effort-related costs (Kennerley et al., 2009; Hillman and Bilkey, 2010). Importantly, Kennerley et al. (2009) also showed that the spike rate of some neurons that respond to predicted reward magnitude is modulated by the anticipated number of lever presses that the monkey would have to make to receive the reward. This suggests that the ACCs may be important for calculating the net value of rewards, i.e. the value of the rewards discounted by the effort required to obtain them. Croxson et al., (2009) tested this hypothesis in humans using fMRI. Subjects performed trials where they would receive either a high reward or a low reward following either a high or low level of effort. They examined activity at the time of cues that instructed subjects as to the level of reward available and how much effort was required for its receipt on each trial. They found that activity in a dorsal portion of ACCs (in posterior MCC/RCZ) showed an interaction between the level of effort and the reward value at the time of the cues that instructed the actions. Thus, it appears that the ACCs may process anticipatory net reward values at the time that actions are instructed.

This thesis examines two main questions (i) does the ACCg process similar information to the ACCs? and (ii) does information processing in this area conform to the principles of RLT? In the previous chapters, activity was found in the ACCg which codes for the erroneous predictions of others, in a manner that conforms to RLT. However, to examine whether this area is engaged when processing predicted values, the other major component of RLT, it is important to examine activity at the time that other's make predictions about the outcomes of their actions. Does the ACCg process the predicted value of effort discounted rewards that are to be received by another? In this study, I examine activity in the ACC at the time of cues that instruct either a first-person or a third-person to perform a series of actions. Specifically examining whether activity in the ACCs varies with the net value of first-person rewards and activity in the ACCg varies with the net value of third-person rewards.

Subjects performed two tasks. For the first task (the effort task) they performed a series of cued actions to receive either a high or low financial reward. The actions entailed the cancellation of a series of targets by making responses on a keypad. The number of cancellations required pertained to the level of effort necessary to receive the reward (2, 3, 8 or 12 cancellations). In the second task, subjects monitored trials where the third-person (confederate) was required to make cancellations. On each trial, the subjects had to indicate the magnitude of the reward that would be received by the third-person (16p, 4p or 0p if the

third-person made the incorrect number of cancellations). The level of effort, reward and who would respond on each trial, was cued by a series of circular stimuli as used in previous studies (Knutson and Cooper, 2005; Croxson et al., 2009; Kurniawan et al., 2010). Activity time-locked to these instruction cues was examined. This allowed for an examination of brain activity that varied with the level of first-person and third-person anticipated rewards, anticipated effort and anticipated net reward value (reward/effort).This experiment tests two hypotheses: Firstly, that activity in the ACCs will vary with the net-reward value on the first-person trials, and secondly, that activity in the ACCg will vary with the net-reward value on the third-person trials.

5.3 Methods

5.3.1 Subjects

Subjects were sixteen, healthy right-handed participants screened for neurological disorders (aged between 18 and 32; 13 female). Two subjects were excluded from the analyses. Both subjects failed to maintain a belief in the deception and one of these subjects failed to perform the judgement task (see below) better than chance (one male). Subjects were paired up with one of two confederate participants, who they believed were also naïve participants. The subjects believed that they would be paid for their participation based on their performance of the task during a scanning session (see below). They also believed that the confederate would be paid based on their performance in the same manner.

5.3.2 Experimental design

The aim of this experiment was to examine the processing of cues that instructed a first-person and a third-person as to how much reward they would receive following the exertion of effort. Subjects performed a task over two days with a training partner (confederate). On the first day, the subject and the confederate learned the associations between a set of instruction cues, a financial reward, and how much effort they were required to expend for its receipt. On the second day, both agents continued to perform effortful actions to receive rewards. During this session, the subject performed these trials whilst inside the MRI scanner, with the training partner situated in the adjacent control room.

A 2x2x2 Factorial design was employed to examine activity time-locked to instruction cues (see fig.5.1). The first factor was Agency. On each trial either the subject (first-person) or the confederate (third-person) performed a series of cued button presses (or 'cancellations') on a keypad to receive a reward. The second factor was the reward level that was obtainable on each trial. This could be either high (HR), if 16p was obtainable on the trial, or low (LR), if only 4p was obtainable. The third factor was the level of effort. There were four levels of effort (2,3,8 or 12 responses), which corresponded to the number of cancellations (cued button presses) that were required to receive the reward. These were collapsed into either low effort

(LE; 2 or 3 cancellations) or high effort (HE; 8 or 12 cancellations) conditions for a factorial design (see below). All cues were colour-coded on each trial, such that the first-person responded when stimuli were blue and the third-person when stimuli were brown. On each trial, the instruction cues signalled the level of reward available and effort required by either the first-person or the third-person. The instruction cue stimuli were based around those used in previous studies (Knutson and Cooper, 2005; Croxson et al., 2009) that investigated first-person reward prediction processing. The stimuli were 80mm diameter circles containing crosshairs. The position of the crosshairs indicated both the amount of reward that was obtainable and the number of cancellations (cued button presses) required to receive that reward. Reward was represented vertically on the circle (16p was high on the circle, 4p was low). Effort was represented horizontally with increasing levels of effort represented from left to right.

In total there were 16 different trial types dependent on the reward level, effort level and the agent performing the cancellations. There were eight different trial types for each level of Agent.

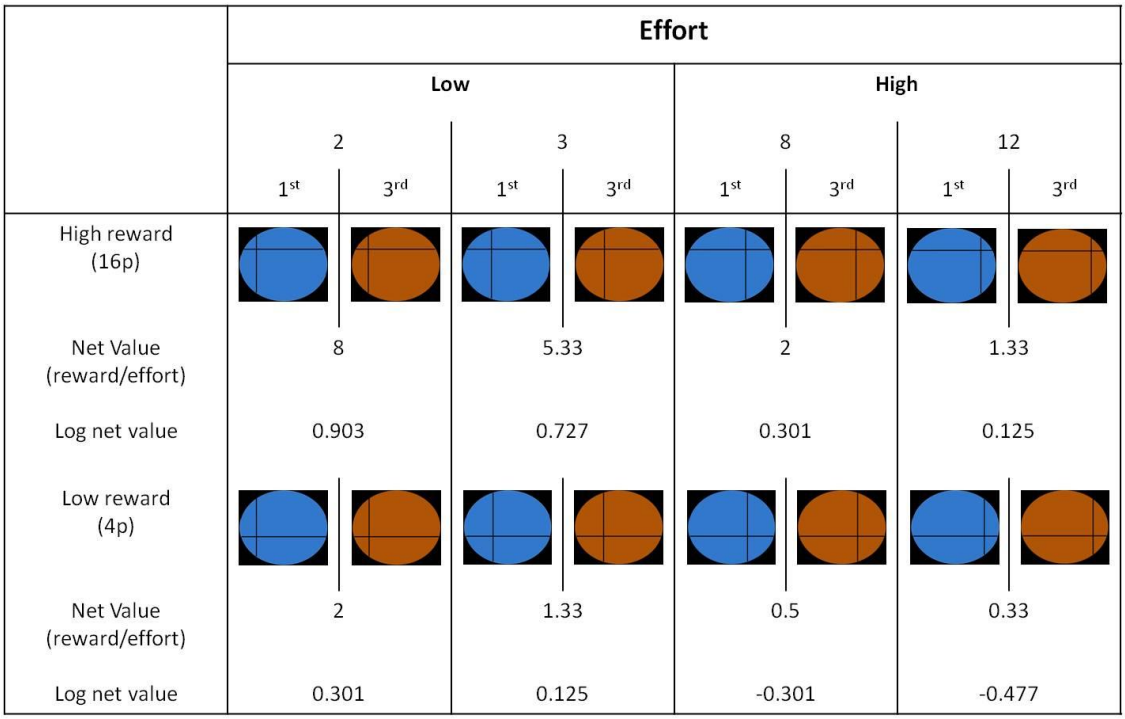

Fig. 5.1 Experimental design. Instruction cue stimuli signalled the level of effort (2, 3, 8, or 12 cancellations) and the level of reward on each trial. The height of the crosshairs was an index of the financial reward (high = high reward, 16p ; low = low reward, 4p). They were also an index of the required level of effort (increasing from left to right). The effort levels were represented on the stimuli on a log-scale, to ensure a clear disparity between the position of the crosshairs for 2 and 3 button presses. The stimuli were also colour-coded to indicate who was required to perform the effort task (either the first-person or the third-person) and the reward that was available. The net reward values reflect the level of reward divided by the level of effort for each stimulus. The log transformed net reward values are shown underneath.

The aim of this study was to investigate whether activity in the ACC signalled the value of a reward discounted by effort. Two strategies were employed to examine the processing of anticipated reward and effort at the time of the instruction cues:

1. Parametric. Although a factorial approach was used, the majority of the results reported in this chapter are from a set of parametric analyses. Given the experimental design, a factorial approach would lack sensitivity to identify activity which responded to net reward values (and log net reward values). Specifically, the large number of conditions used in the experiment (16 conditions in total) resulted in a small number of repetitions (12 of each instruction cue) of each type of cue. This was a necessity to keep the duration of the experiment practical (the scan was 48.5minutes in duration). A parametric approach, examining activity which varied with the net level of reward or the level of effort, therefore had greater sensitivity when compared to the factorial approach:
    a. *Net Reward Value* - To examine whether activity in any voxels varied with the reward values discounted by effort, a net reward parameter was created for the net reward values (reward/effort; see fig. 5.1 for table of net reward values). Separate parameters which were scaled by the net reward value were created for the first-person and third-person trials. To identify voxels which responded exclusively to the net reward values at the time of the third-person instruction cues, voxels in which the timecourse varied significantly (p<0.05unc) with the first-person net reward values were excluded. To identify voxels which responded exclusively to first-person net reward values, voxels in which the timecourse varied significantly (p<0.05unc) with third-person net reward values were excluded. To examine areas in which activity was scaled by net reward values regardless of the Agent who was to receive it, a conjunction was conducted between the two parameters.
    b. *Log Net Reward Value* – There is evidence that animals and humans represent effort related costs on a logarithmic scale (Brunner et al., 1992; Croxson et al., 2009; Kurniawan et al., 2010). In addition, Croxson et al., (2009) found that activity in the ACCs varied with the log-scaled value of rewards. Thus, in this study, the net reward values were log-transformed. The log-transformed values (see fig. 5.1) were then used as the parameters in a new set of analyses, which used the same approach as outlined in (a) above.

c. *Effort* - The final parametric analysis examined whether activity in any voxel was scaled with the level of effort. One parameter was created for the level of effort on the first-person trials and one parameter for the level of effort on the third-person trials. As outlined in (a) for net reward value, to examine activity that responded exclusively to effort for an Agent, voxels in which activity varied with the anticipated level effort for the other Agent were excluded.

2. Factorial analysis. In addition to the parametric analyses, a factorial design was employed for two reasons: Firstly, to examine undiscounted reward values. As there were only two levels of reward (16p or 4p), there was not enough variation for the values of reward to be parameterised. However, undiscounted rewards could be examined in a factorial design. Secondly, this design had the added benefit of allowing PSTH plots to be created for the peak voxels identified by the parametric analyses (these are included in appendix D). The 2x2x2 factorial design outlined above was used to examine activity time-locked to the instruction cues. To identify voxels which processed reward value (a differential response between HR and LR, regardless of the level of effort) a contrast was conducted to look for the main effect of reward, regardless of Agent or Effort. To identify voxels which responded exclusively to first-person reward value, regardless of the level of effort, a two-way interaction was conducted between Reward and Agent. Voxels which showed a main effect of third-person reward were excluded (p<0.05unc). To identify voxels which responded exclusively to third-person reward value, regardless of the level of effort, a two-way interaction was conducted between Reward and Agent. Voxels which showed a main effect of first-person reward were excluded (p<0.05unc).

5.3.3 Trial Structure

Each trial (see fig.5.2) began with one of 16 different colour-coded instruction cues. These cues indicated both the level of reward that was available on each trial and also the level of effort required for its receipt. The colour of these cues also indicated who would have to perform the cancellations on each trial (blue for the first-person, brown for the third-person). Following the instruction cue there was an effort period, during which subjects made cancellations by pressing buttons on a keypad. During the effort period on the first-person trials, subjects were

required to make a series of cued button presses (cancellations). On the third-person's trials the cancellations were actually pre-programmed computer controlled responses (see below for more details). At the end of the effort period, a stimulus then displayed the number of cancellations that had been made during the effort period. Following this stimulus, a trigger cue (three lines, with 16p over the left hand line, 4p over the middle line and 0p over the right hand line) was presented on the screen, which cued the first-person or the third-person to make a judgement of the amount of reward that would be received by the other agent on that trial. Each line corresponded to one line on the keypad. Subjects had 750ms to make their response. If they did not respond in this time window it was classified as an incorrect response. Following this, a feedback cue indicated the accuracy of the judgement ("correct" if the judgement was correct and "-10p" if incorrect) and then finally a feedback cue indicated the reward received by the agent who performed the cancellations (16p, 4p or 0p). The full trial structure and timings can be seen in figure 5.2.In total there were 192 trials, 96 first-person trials where the subject made cancellations and 96 where they monitored the third-person's cancellations. Each instruction cue was presented on twelve trials.

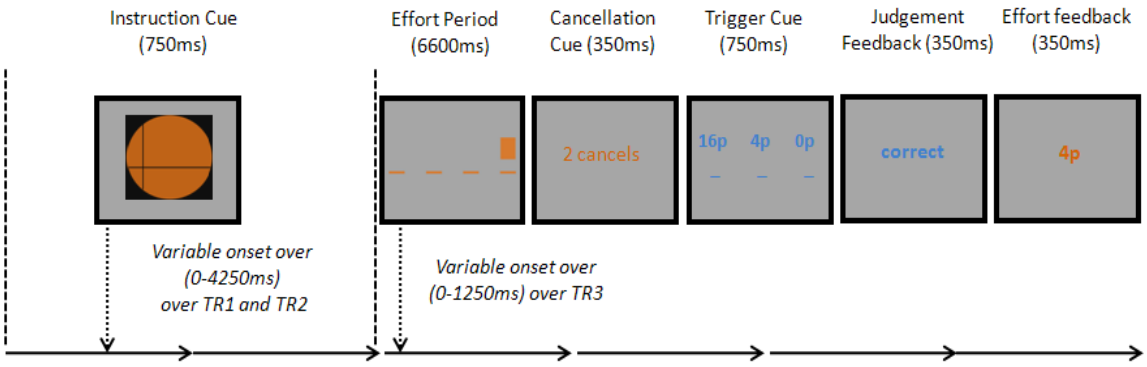

Fig. 5.2.Trial Structure. Each trial began with one of 16 different colour-coded instruction cues. If the cue was blue, the first-person would perform a series of cancellations and the third-person would make a judgement of the amount of reward the first-person would receive; If brown, a third-person would perform a series of cancellations and the first-person would make a judgement of the reward to be received by the first-person. The instruction cue onsets were uniformly and randomly jittered over the first two TRs. Following the instruction cue there was an effort period. During this time a series of cued button responses were made on the keypad. The responses were cued by the position of a square target over one of four lines which corresponded to the four buttons on the keypad. Following the effort period, the number of cancellations made (i.e. target button responses) was presented on the screen. After this, a trigger cue was presented, instructing a judgement to made by one agent about the reward to receive by the other agent. Feedback then indicated the accuracy of that judgement, and finally feedback was presented for the outcome of the performance of the effort period.

5.3.4 Task

Subjects performed two tasks during scanning. On first-person trials subjects performed an 'effort task' where cancellations were made in order to receive financial rewards. On the third-person's trials subjects performed a judgement task, monitoring the third-person's performance of the effort task and indicating the amount of reward they would receive.

5.3.4.1 Effort Task

During scanning, subjects performed trials where they were required to make the correct number of cancellations to receive a financial reward. The 'effort task' required subjects to make a series of cancellations during the effort period. Each cancellation was cued by the position of a square stimulus above one of four lines on the screen. Each line on this cue, corresponded to one button on the keypad. The position of the square highlighted a target button. A cancellation constituted one press of the target button i.e. one finger movement of one finger on the right hand. Once this target button was pressed, the position of the square would move to highlight a new target button. Each target button was always different to the previous. Subjects could make up to 14 of these cancellations during the fixed time window of the effort period (6600ms). Target buttons were randomised across the experiment and within each effort period. Subjects were therefore unable to make any prediction about which button would be the next target.

During scanning, subjects were told that they were accumulating monetary rewards for their performance on this task. As such subjects believed that they were earning the reward available on each first-person trial, if they performed the correct number of cancellations. Subjects were told that if they performed every cancellation correctly they would accumulate £10 as payment for the experiment. However, unbeknown to the subjects, they would be paid £10 for participation regardless of their task performance.

5.3.4.2 Judgement Task

In addition to the effort task, subjects also performed a judgement task on the trials where the third-person was performing the effort task. For this task the subjects were required to indicate the level of reward that would be received by the third-person, which could be 16p or 4p for the correct number of cancellations or 0p if the number of cancellations was incorrect.

Subjects were required to perform this judgement on every trial performed by the third-person. Subjects believed that they were punished for each incorrect judgement (when the subject indicated that the amount of money earned by the confederate on the third-person trials, was different from the amount they would actually earn) by 10p being removed from the money they were accruing on the effort task. A correct judgement left the rewards accumulated during the effort task the same. Thus, if subjects performed every set of cancellations correctly, but every judgement incorrectly, they would receive no payment for the experiment. Thus, subjects were motivated to perform both tasks to the same degree of accuracy. The punishment therefore ensured that subjects attended to the rewarding value and the effort information contained in the instruction cues on the third-person's trials. Importantly, the punishment used as the motivation for the subject on the third-person's trials, was unrelated to the anticipated reward and effort level that would be processed at the time of the instruction cues.

The effort task used in this study is very similar to that employed by Croxson et al., (2009) that was used to investigate first-person effort-discounting. Given this similarity, it is important to note the differences between the task employed here and that used by Croxson et al., (2009). Croxson et al., (2009) used 8 different instruction cues in which crosshairs indicated the level of effort and the amount of reward obtainable. There are two important differences between the effort task employed in that study and that employed here. Firstly, unlike in this study, there were no constraints placed upon the time which subjects had to make cancellations. Secondly, unlike in this experiment, subjects were only presented with the correct number of targets to be cancelled. However, these two aspects of their task were not suitable for the purposes of this study. Crucial to the design of this study, was that subjects were required to make a judgement on the reward to be received by a third-person. This task ensured that subjects attended to the effort and reward levels at the time of the instruction cues on the third-person's trials. Without a temporal constraint on the effort period, there would be no possibility of making an incorrect number of cancellations, and thus, the confederate would not make errors on the effort task. Without confederate errors, the subject could perform the judgement task by attending to the level of reward at the time of the instruction cues on the third-person's trials and not the level of effort. Thus, in this experiment, a temporal window was a necessity. In addition, in this experiment subjects could cancel up to 14 targets, more than the maximum instructed number of 12, regardless of how many cancellations they were required to make. This created the potential for catch trials, where the confederate made an error in the number of cancellations, which the subject would need to identify correctly in

order to maximise their own financial rewards. These two distinctions from the task used by Croxson et al., (2009) required subjects to attend to both the effort and reward level on every confederate trial.

5.3.5 Procedure

5.3.5.1 Training Session

Subjects were trained in two phases one day prior to scanning. In the first phase, the subject was seated in front of a monitor with a confederate (third-person). They were each provided with a response keypad. A series of visual stimuli were presented on the screen. Both the confederate and the subject performed the effort task on separate trials. During this session both the confederate and subject learned the contingency between the position of the crosshairs on the instruction cue stimulus, the amount of reward (16p or 4p) and the required number of cancellations (button presses) to receive the reward. They were informed before this that there would be two levels of reward and that they would have to make two, three, eight or twelve button presses. During training there were 64 'first-person' trials where the subject performed the cancellations and 64 'third-person' trials performed by the confederate. The subjects were told that the rewards were fictional during training and their payment for the experiment would be based solely on performance during the scanning session.

In this session, as the subjects were seated next to the confederate, the confederates' performed the effort task on separate trails from the subject. As the confederates were paired with multiple different subjects throughout the piloting and experimental phases, they were highly over-trained on the effort task. To ensure that subjects maintained the belief that the confederates' were naïve participants like themselves, they were told to make deliberate errors in the number of cancellations performed during the first phase of training, to mimic the learning of a real participant.

In the second phase of training, subjects practised the task that would be performed during the scanning session (see below). The subject performed this from inside a mock scanner, with the confederate seated in front of a monitor adjacent to the mock scanner. The subject was played the sound of a genuine scanner's EPI sequence via headphones. During this training phase and during scanning the responses on the third-person trials were computer controlled.

5.3.5.2 Scanning Session

Prior to scanning subjects were shown the confederate seated in front of a monitor in the control room next to the scanner. They were told that they would see all of the responses of the third-person in real-time inside the scanner. In fact these responses were all computer controlled, pre-programmed responses. The apparent reaction times of the confederate during the effort task were pseudorandomly organised. The reaction times of the second to twelfth button presses fitted a normal distribution around a mean (525ms), with a range of 325ms to 725ms. The confederate's reaction times to the first target were extended, to reflect the unpredictability of the onset of this target. These formed a normal distribution around a mean of 600ms, with a range of 400ms to 800ms. These timings were based on the reaction times of five participants during a pilot experiment. Key for the design of this experiment was that subjects attended to both their own instruction cues and those of the third-person in the same manner. There was one potential caveat to the judgement task used to motivate subjects to attend to the instruction cues of confederate. Specifically, if the confederate performed the correct number of cancellations on every trial, the subject could, over time, learn to perform the judgement without attending to the level of effort, only the reward level. To address this potential confound, errors were pre-programmed into the behaviour of the confederate. On nine of the trials where the effort task was performed by the confederate, the number of cancellations performed was not correct for the instruction cue presented. These 'catch' trials were used as an index of the extent to which subjects were attending to the effort expended by the confederate.

5.3.6 Experimental timing

An important feature of this study was that activity was examined that was time-locked specifically to the instruction cues. These cues signalled the anticipated reward and anticipated effort on both the first-person and the third-person trials. To examine activity time-locked to the instruction cues, a variable delay was introduced between the instruction cue and the onset of the effort period, and also the offset of the effort feedback. This allowed for BOLD activity time-locked to the instruction cues to be isolated, without the contaminating effects of either prior or subsequent trial events. Events in each trial took place across six TRs (0s–15 s; TR=2.5 s). To optimally sample the instruction cue related activity, a randomly varying interval

between the scan onset and these cues was introduced over the first and second TRs. This achieved an effective temporal sampling resolution much finer than one TR for the conditions of interest. These intervals were uniformly distributed for each of the 16 different conditions, ensuring that Evoked Haemodynamic Responses (EHRs) time-locked to the instruction cues were sampled evenly across the time period following each instruction cue. The effort period and all events that followed were jittered over the first 1250ms of the third TR.

5.3.7 Functional imaging and analysis

5.3.7.1 Data acquisition

T1-weighted structural images were acquired at a resolution of 1×1×1 mm using an MPRAGE sequence. 1164 EPI scans were acquired from each participant. 34 slices were acquired in an ascending manner, at an oblique angle (≈30˚) to the AC-PC line to decrease the impact of susceptibility artefact in subgenual cortex (Deichmann et al., 2003). A voxel size of 3×3×3 mm (25% slice gap, 0.8 mm) was used; TR=2.5s, TE=32, flip angle=81°. The functional sequence lasted 48.5 minutes. Immediately following the functional sequence, phase and magnitude maps were collected using a GRE field map sequence (TE_1 = 5.19ms, TE_2 = 7.65ms).

5.3.7.2 Image preprocessing

Scans were pre-processed using SPM8 (www.fil.ion.ucl.ac.uk/spm). The EPI images from each subject were corrected for distortions caused by susceptibility-induced field inhomogeneities using the FieldMap toolbox (Andersson et al., 2001). This approach corrects for both static distortions and changes in these distortions attributable to head motion (Hutton et al., 2002). The static distortions were calculated using the phase and magnitude field maps acquired after the EPI sequence. The EPI images were then realigned, and coregistered to the subject's own anatomical image. The structural image was processed using a unified segmentation procedure combining segmentation, bias correction, and spatial normalization to the MNI template (Ashburner and Friston, 2005); the same normalization parameters were then used to normalize the EPI images. Lastly, a Gaussian kernel of 8 mm FWHM was applied to spatially smooth the images in order to conform to the assumptions of the GLM implemented in SPM8.

5.3.8 Statistical Analysis

5.3.8.1 First-Level analyses

First-level GLMs were created for both factorial and parametric analyses.

1. *Factorial Analysis.* There were 10 event types. Each event-type was used to construct a regressor by convolving the stimulus timings with the canonical HRF. Each of the eight conditions was modelled as a separate regressor. In addition, one regressor modelled the activity during the effort periods (regardless of whether it was a first-person or third-person trial) and another regressor modelled the onsets of the other trial elements on every trial. Trials in which the subject failed to perform the correct number of cancellations during the effort period, failed to respond within 750ms of the onset of the trigger cue for the judgement task, or failed to make the correct response on the judgement task were modelled separately as an extra regressor. This regressor included the onsets from all of the trial elements from these events. The residual effects of head motion were modelled as covariates of no interest in the analysis by including the six head motion parameters estimated during realignment.

2. *Parametric Analysis.* Five GLMs were created at the first-level which employed a parametric approach. Each of these GLMs was constructed using the same events as those used in the factorial analysis (see above). For these GLMs, however, the instruction cue regressors were collapsed down into one regressor for the first-person instruction cues and one regressor for the third-person instruction cues. The parameters outlined above (section 5.3.2) were then used as first-order parametric modulators of first-person and third-person instruction cue events. To examine activity which varied with the net reward value parameter, the net reward parameters were orthogonalised with respect to the effort parameters. This ensured that activity which varied with either first-person net reward values, third-person net reward values or both could not be explained by the level of effort alone. The second GLM was identical, except that the log-transformed net reward values were used for the values of the parametric modulator. The third and fourth GLMs were similar except that the effort parameters were orthogonalised with respect to the net value and log net value parameters. The final GLM contained both the log net reward value and net reward

value parameters for both the first-person and third-person trials without orthogonalisation. This allowed for a contrast to be conducted to determine whether the net or log net reward value parameters were a better fit of the data.

5.3.8.2 Second-level analysis

Random effects analyses (Full-Factorial ANOVAs) were applied to determine voxels significantly different at the group level. SPM{t} contrast images from all subjects at the first-level were input into second-level full factorial design matrices. *F*-contrasts were conducted in each of the second-level Random-effects analyses. For the whole brain analyses, FDR correction was applied. To test the specific hypothesis of the thesis, 80% probability masks of the ACCg and ACCs were created and used as the search volumes for small volume correction. In addition, given the similarity in design between this study and that of Croxson et al., (2009), small volume corrections were applied as a sphere with 8mm radius around the peak coordinates from their analyses. This correction was applied by making a mask combining each of the spheres around the peak coordinates (The coordinates are included in Appendix E).

5.4 Results

5.4.1 Behavioural results

The subjects performed two tasks whilst inside the MRI scanner. They performed a first-person effort task and also a judgement task on the trials where the third-person was performing the effort task. For the effort task, subjects were required to make either 2, 3, 8 or 12 cancellations in order to receive a financial reward. An important issue in this experiment was that the cancellation task constituted effort and not difficulty. In previous studies, the effort period (Botvinick et al., 2009; Croxson et al., 2009) was not constrained by a time-period, and as such cancelling a large number of targets was not more difficult than a small number. In this chapter, the fixed response window (6600ms) may have caused subjects to find it more difficult to complete the 12 cancellations than the 8, 3 or 2 cancellations. This would confound any interpretation of the results as being effort-related, as effort-related activity would have been confounded with activity occurring due to the difficulty of the task. To determine whether effort was confounded with task difficulty, the behavioural accuracy of subjects across each of the four effort levels were examined on the first-person trials (see fig.5.3). A repeated measures ANOVA was performed examining the effect of effort on task accuracy (% of correct responses). No main effect of effort on task accuracy was identified ($F(2.18,28.37) = 2.098$, $p>0.1$). Planned pairwise comparisons between levels of effort revealed no significant differences in accuracy between any two levels of effort for 2<>3 (t(13) = 1.528, $p= 0.151$); 2<>8 (t(13) = 1.749, $p= 0.104$); 2<>12 (t(13) = 0.366, $p= 0.720$); 3<>12 (t(13) = 1.046, $p= 0.315$); 8<>12 (t(13) = 1.249, $p= 0.234$)) apart from a significantly lower accuracy for 8 cancellations than 3 cancellations ($t(14)= 2.621, p<0.05$). Fig.5.2 shows the means and standard error for task accuracy at each level of effort. Importantly, the trials where 12 cancellations were required were performed at the same level of accuracy as trials where 2, 3 or 8 cancellations were required, suggesting that having a fixed response window did not impact upon the difficulty of the effort task.

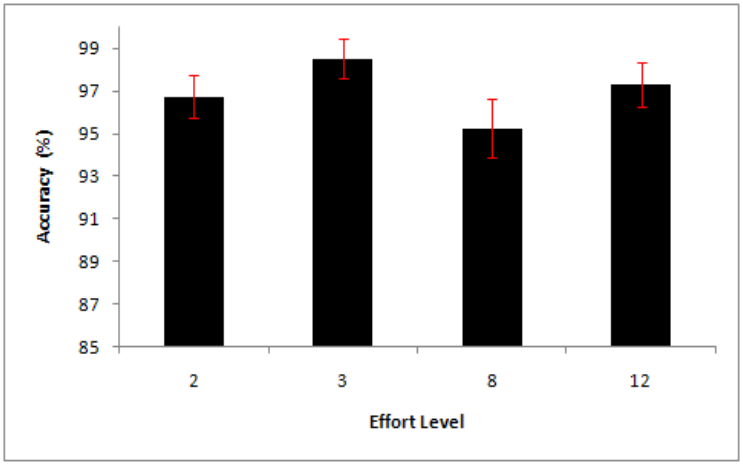

Fig. 5.3 First-person effort task accuracy

The second task performed by subjects was a judgement of the reward that would be received by the third-person. Subjects were required to monitor the responses of the third-person and indicate whether they would receive the high reward (16p), low reward (4p) or no reward (0p) on each trial. Performance on this task was an important index of subjects' understanding of the level of reward available and the effort necessary for its receipt, on the third-person's trials. It was of particular importance that subjects performed 'catch trials', where the third-person made the incorrect number of cancellations, above the chance level of accuracy (33.3%). On these trials subjects could not perform the judgement task correctly without attending to the reward level and required effort level at the time of the instruction cue and also the number of cancellations actually made by the third-person (see fig.5.4 for results). A paired samples *t*-test revealed that subjects' overall task accuracy (mean = 93.93%) was significantly better than chance ($t(14) = 54.5$, $P<0.0001$). On the catch trials the accuracy was lower (mean = 78.64%) but still significantly greater than chance ($t(14) = 12.76$. $P<0.001$). These results indicate that subjects were attending to the reward value, the level of effort and the number of cancellations made by the third-person. A repeated measures ANOVA also revealed that there was no effect of the level of effort on task accuracy ($F(2.549, 33.142) = 1.97$, $p = 0.145$). Planned pairwise comparisons between 2<>3 ($t(13) = 0.798$, $p= 0.439$); 2<>8 ($t(13) = 0.332$, $p= 0.745$); 2<>12 ($t(13) = 20.32$, $p= 0.061$); 3<>8 ($t(13) = 0.660$, $p= 0.520$); 3<>12 ($t(13) = 1.290$, $p= 0.220$); 8<>12 ($t(13) = 2.034$, $p= 0.063$)) revealed no significant differences in task accuracy between any two levels of effort. The accuracy of subjects on the judgement task therefore cannot be explained by the level of effort that the third-person was required to exert. The performance of subjects on both tasks indicates that they were processing the reward value and the effort level at the time of the instructions cues on both first-person and third-person trials.

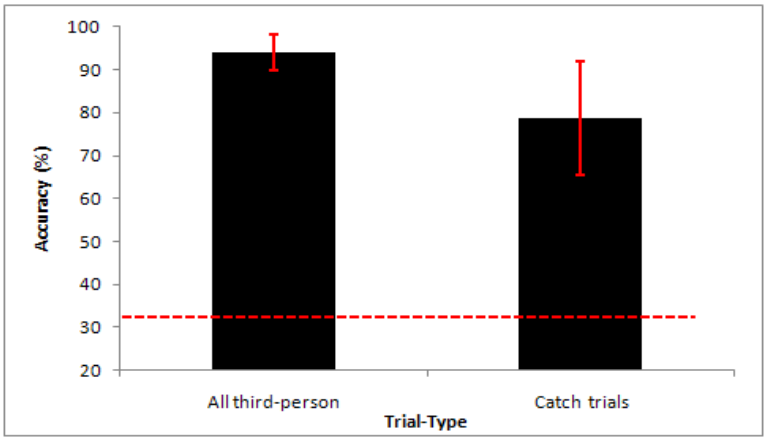

Fig. 5.4. Judgement task accuracy of the first-person. Subjects performed a judgement of the level of reward that would be received by the third-person on each trial. Catch trials were trials where the third-person performed the incorrect number of cancellations. 'All third-person' shows the accuracy across both catch (0p), 4p and 16p trials. The dotted line represents the chance level on the task.

5.4.2 Imaging results

5.4.2.1 Net Reward Value

A parametric analysis was conducted to examine brain activity which was scaled parametrically with the net value of rewards discounted by effort. Two different parameters were used, a net reward value parameter (reward/effort) and a log-transformed net reward value parameter. These were fitted to the instruction cues of the first-person and the third-person trials and the different parameters were entered into separate GLMs. Two compare the two parameters in terms of their fit to the data, they were entered into a separate GLM, where they were not orthogonalised with respect to each other.

To examine activity which varied with the net reward value on the first-persons trials an *F*-contrast was applied on the parameter, which had been orthogonalised with respect to the effort parameter. To examine voxels which varied exclusively with first-person net reward value, voxels which responded to third-person net reward value were excluded. Activity was found in the ventral striatum, putatively in the nucleus accumbens (- 8, 14, -4, t(104) = 2.92, Z = 2.8; $p<0.05$ svc around the peak coordinate from Croxson et al. 2009) that varied with the log transformed net value on the first-person trials (see fig. 5.5). Activity in this area was explained marginally better by the log-transformed parameter (t(104), $p<0.05$ unc) than the untransformed net values. This area also showed a marginally significant interaction between first-person Reward and Effort in the factorial analysis (F(1,130) = 5.55, Z = 2.06, $p<0.05$ svc).

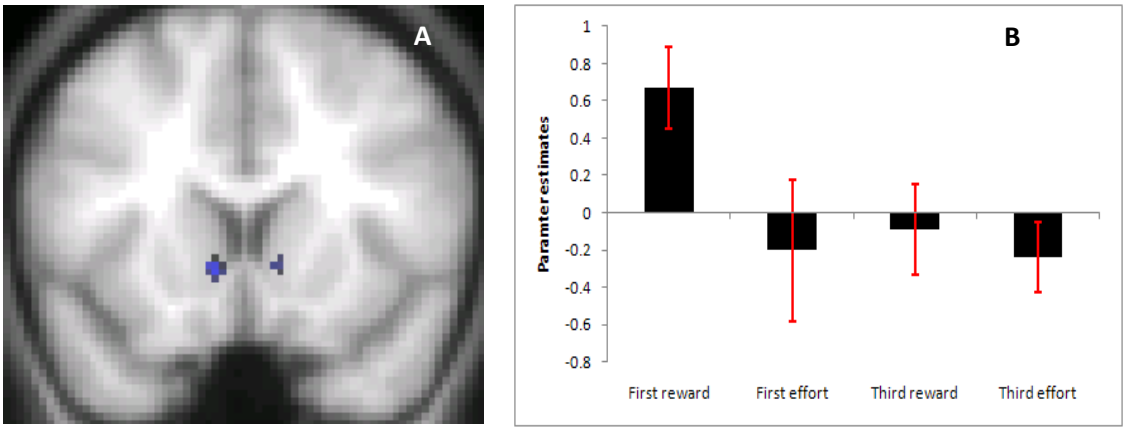

Fig. 5.5. First-person log net reward values. (A) Activity shown in the ventral striatum (-8, 14, -4, Z = 2.8, P<0.05svc), (B) Beta coefficients of peak voxel for the log net reward values for the First-person, Third-person log net reward and effort parameters. The beta values are taken from the parametric analysis in which the reward parameters were orthogonal to the effort parameters.

To examine whether activity in any brain area was scaled by the net reward value on the third-persons trials, an *F*-contrast was applied to the third-person net (or log net) value parametric regressors. To ensure that any area that was identified by this contrast was sensitive exclusively to third-person net reward value, voxels which showed a significant effect (P<0.05unc) of the first-person instruction cue events, the first-person effort parameter or the first-person net reward parameter were excluded. Activity was found to vary with the log net reward parameter in only one area, the ACCg (4, 22, 20,F(1,91) = 9.8 Z = 2.8, *p*<0.01svc; putatively in MCC area 24b', see fig.5.6). Activity in this area was explained marginally better by the log-transformed net value parameter than the untransformed net reward values (F(1,104) = 2.7, Z = 2.3, *p*<0.05unc). This area also showed a significant interaction between Effort and Reward in the factorial analysis (F(1,130) = 11.34, Z = 3.09, *p*<0.05svc). No voxels were found in the ACCs that responded exclusively to third-person net -value. No voxels within the reported cluster encroached into the ACCs and did not overlap with any other cluster reported as activated in the ACC in any contrast.

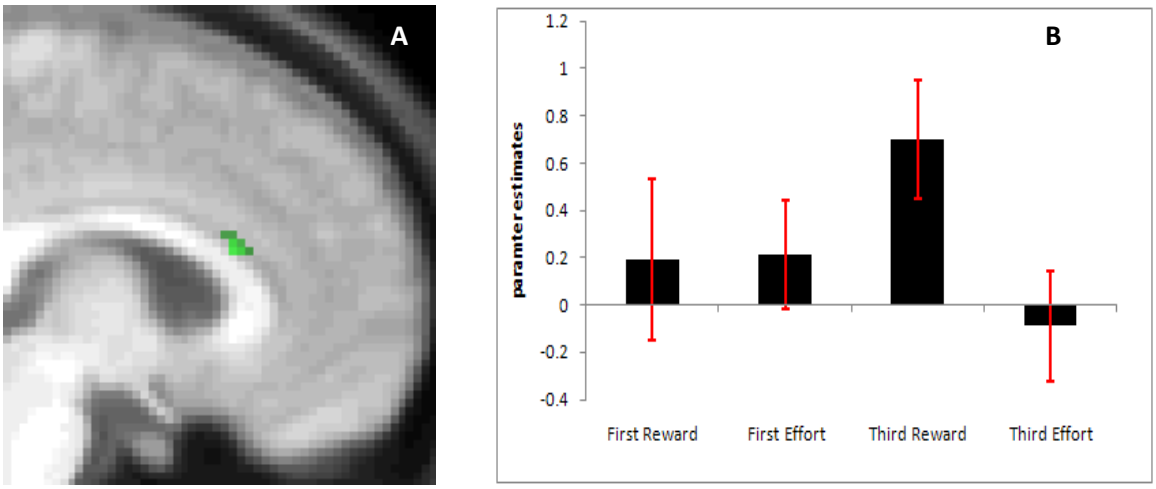

Fig. 5.6. Third-person log net reward value related activity. (A) Activity shown in the gyral portion MCC area 24a'/24b' (4, 22 ,20, Z = 4.2; P<0.05svc), (B) Plots of the Beta coefficients from the parametric analysis. The beta values are taken from the parametric analysis in which the reward parameters were orthogonal to the effort parameters.

5.4.2.2 Reward Value

To examine whether any brain area responded to the magnitude of the reward not discounted by effort, main effect contrasts were performed in the factorial analysis. It was not possible to analyse reward values parametrically as there were only two levels of reward (high or low). There was therefore not enough variation in the reward values for them to be parameterised. An *F*-contrast comparing HR with LR conditions for the first-person instruction cues revealed no voxels which responded to first-person undiscounted reward. An F-contrast comparing HR with LR at the time of the third-person instruction cues revealed activity in four areas that showed a main effect of reward. This included a region extending across gyrus rectus, the superorbital gyrus and the medial superior frontal gyrus commonly referred to as Ventromedial Prefrontal Cortex (VmPFC; 6, 42, -4, F(1,130) = 18.50, Z = 3.9, p<0.05FDR; extending across areas 32, 10p and 10r (Ongur and Price, 2000)), a frontal polar region in area 10 (10, 64, 8; F(1,130) = 18.52, Z = 3.9, p<0.05FDR), and also an activation that extended across a large portion of the gyral MCC (area 24a'/24b'; 4, 16, 26; F(1,130) = 14.15, p<0.05FDR, Z = 3.66). Importantly, this portion of the ACCg did not overlap with the cluster in which activity varied with the log net reward value reported above (see fig.5.7). Examination of the PSTH shows that this area made significantly negative responses to undiscounted third-persons rewards, with the most negative response to the highest reward.

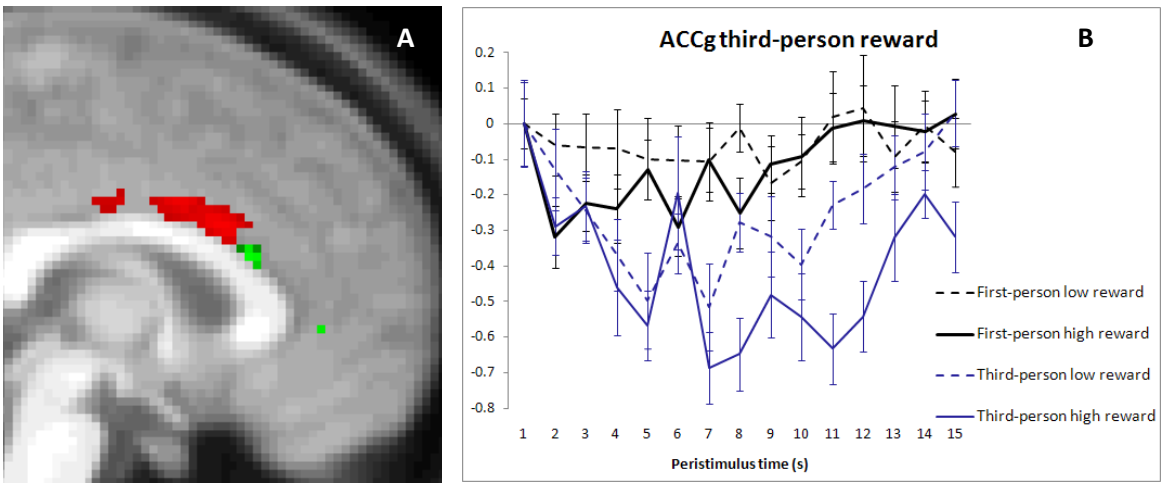

Fig. 5.7. Third-person undiscounted rewards. (A) Activity shown in the gyral portion of midcingulate area 24b' in red which showed a significant difference between HR and LR at the time of the third-person instruction cues (4,16,26; p<0.05FDR). This cluster was spatially separated from the region which showed a parametric response to the third-person log net reward values, shown in green (B) PSTH of the responses in the peak ACCg voxel. For display purposes the conditions are collapsed into first-person and third-person, HR and LR conditions.

5.4.2.3 Effort

To examine activity that varied with the level of effort, separate *F*- contrasts were applied to the parametric regressors which varied with the level of effort for either the first-person or third-person. Activity was found to vary with the first-person effort level in several brain areas (table 5.1.). Activity which varied exclusively with the third-person effort level was found in the lingual gyrus (p<0.05FDR). A conjunction across both first-person and third-person effort parameters showed that activity in a posterior portion of the sulcal MCC/RCZ (0, -22, 52; F(1,104) = 16, Z = 3.44, p<0.05svc, putatively area 24c' see fig.5.8) varied with the level of effort, regardless of whether the effortful actions were to be performed by the first or third-person. This was the only region to show such a profile. This region also showed a significant main effect of effort, regardless of the level of Agent, in the factorial analysis (p<0.05unc). This area did not overlap with the activations reported in the ACCg.

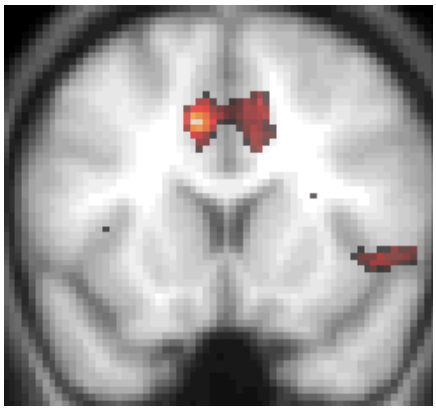

Fig. 5.8. Conjunction between first-person and third-person effort parameters. (A) Activity shown in the posterior portions of the sulcal midcingulate area 24c', (B) PSTH of the peak voxel from the conjunction between first and third-person effort. For display purposes PSTH values were collapsed across the two levels of Reward and Agency, to give the main effect PSTH for High and Low Effort.

Anatomical region	***MNI Coordinate***	***Z-value***
Parietal		
Intraparietal sulcus	52, -28, 50	3.2
Cingulate		
Anterior Cingulate Sulcus	-12, -2, 40	3.19
Frontal		
Middle frontal Gyrus	20, 40, 22	3.38
Temporal		
Posterior Superior Temporal Sulcus	-64, -56, 20	3.97
Parahippocampal gyrus	34, -34, -24	3.51
Central Insular Sulcus	50, 6, -6	3.53
Cerebellum		
Lobule VI	16, -46, 16	4.06

Table 5.1 First-person effort responses. All regions corrected for multiple comparisons (P<0.05FDR) with (1,104) degrees of freedom.

5.5 Discussion

This study tested two hypotheses about the processing of anticipatory reward signals when effortful actions are instructed. The first hypothesis was that activity in the ACCs would be scaled with net reward value (the value of a reward discounted by the amount of effort to be expended) at the time of the instruction cue on the first-person trials. The second hypothesis was that the ACCg would signal this information at the time of the instruction cues on the third-persons trials. In contrast to the first hypothesis activity in the ACCs was not scaled by net reward value and in fact, was not modulated by reward at all. However, activity in this area was scaled with the anticipated level of effort regardless of whether the actions were to be performed by the first-person or the third-person. Activity was found to vary with first-person net reward value in the ventral striatum, in line with previous studies (Botvinick et al., 2009; Croxson et al., 2009; Kurniawan et al., 2010). In support of the second hypothesis, this study showed that activity in a portion of the ACCg, in the midcingulate area 24b', scaled with the log net value of effort discounted rewards when the actions were to be performed by another. In addition, a more posterior portion of the midcingulate area 24a'/b' signalled the anticipated reward, regardless of the level of effort, when this reward was to be received by another. The results therefore support the notion that the ACCg signals the anticipated value of rewards that are to be received by another. Different portions of this area code for the undiscounted magnitude and the effort discounted net value of rewards.

5.5.1 First-person effort based valuation

This study highlighted that activity in the ACCs and the ventral striatum (putatively in the nucleus accumbens) varies with the level of effort and net value respectively. These results therefore suggest that the ventral striatum and the ACCs may form part of a network that is engaged in valuing actions, based on the cost associated with their performance. This notion is supported by anatomical evidence. In non-human primates there is evidence to suggest that the posterior portions of MCC (area 24c') project primarily to dorsal portions of the striatum, with additional weaker connections to the core of the nucleus accumbens (Kunishio and Haber, 1994). In humans diffusion weighted imaging data supports the notion that these two areas are connected, with tractography revealing the presence of white matter fibres between the posterior MCC and a portion of the ventral striatum, including the nucleus accumbens (Beckmann et al., 2009). Both areas 24c' in the ACCs and several portions of the ventral

striatum also show strong connections to the Posterior Cingulate Cortex (PCC) (Pandya et al., 1981), portions of the intraparietal sulcus (Yeterian and Pandya, 1991, 1993; Leichnetz, 2001) and the OFC (Morecraft et al., 1992; Carmichael and Price, 1995; Haber et al., 1995; Cavada et al., 2000) which are known to contain neurons which are sensitive to rewarding stimuli. The portion of the ACCs activated was in the Rostral Cingulate Zone (RCZ), the human homologue of the CMAs. This area has strong connections to areas within the motor system and also with the spinal cord (Picard and Strick, 1996). The connectional anatomy of these two regions makes them well placed to process reward-related and motor related information respectively. Thus, these are highlighted as candidates for processing the value of a reward discounted by the number of actions that are required for its receipt.

There is evidence to support the notion that information is exchanged between the ACC and the basal ganglia during effortful decision-making. Consistent with this interpretation, the region of the ACCs that was activated in this study (in the posterior MCC/RCZ) is connected, although only weakly, to the core of the nucleus accumbens (Kunishio and Haber, 1994), which in turn projects to the ventral pallidum (Spooren et al., 1996). This output nucleus also projects to posterior area 24c'in the ACCs, via the thalamus (Alexander et al., 1986; Alexander and Crutcher, 1990; Groenewegen et al., 1993; Middleton and Strick, 2000; Nakano et al., 2000; Haber and Knutson, 2010). Intriguingly, disruptions to the striato-pallidal connection (Mingote et al., 2008) and also to the striato-cingulate connection (Hauber and Sommer, 2009) perturb normal behavioural patterns on tasks that require a choice between options that have different associated costs. Thus, there is evidence to suggest that the portions of the ACCs and the ventral striatum that responded during first-person effort-based discounting in this experiment are anatomically connected. Disruptions to any of the connections in the loop between them disrupts behaviours that require effort to be evaluated, highlighting the importance of this loop.

Although most of the work in this thesis examines decision-making processes, in this chapter subjects were not presented with a choice. It could therefore be argued that the results of this study do not reflect the processes that underpin decisions between effortful actions, only those that signal the value of instructed actions. However, there is considerable evidence in rats and nonhuman primates that the ACCs and the ventral striatum are involved in the process of guiding choices between effortful options. Lesions to dorsal portions of the ACC (Rudebeck et al., 2006b; Walton et al., 2006; Floresco and Ghods-Sharifi, 2007) and depletions of dopamine within the ventral striatum (Ishiwari et al., 2004; Denk et al., 2005; Salamone et al., 2007; Floresco et al., 2008), particularly to the core of the nucleus accumbens, impair

effort-related decision-making processes. Such lesions bias choices towards less effortful but less rewarding options. Thus, the integrity of both of these regions is required for the process of deciding between courses of actions which require different amounts of effort to be expended. This would suggest that the ACCs and the ventral striatum form an important circuit for valuing different effortful options during decision-making.

Neurophysiological evidence also implicates the ACCs as a region which is engaged in processing the costs associated with a reward. Several studies have found neurons that are sensitive to the value of a number of variables that guide choice behaviour, such as reward probability, reward magnitude and also effort (Amiez et al., 2005; Sallet et al., 2007; Quilodran et al., 2008; Kennerley et al., 2009; Hayden and Platt, 2010). Kennerley et al., (2009) provided convincing evidence that the ACCs of the monkey contains neurons which decrease their firing rate as the level of effort to be expended increases. This would support the result identified in this chapter, where activity in the ACCs was modulated and varied by the level of effort.

Kennerley et al., (2009) also showed that within the ACCs there are neurons in which the coding is specific to one decision variable and also neurons which multiplex this information, coding for interactions between variables. This is intriguing, as Croxson et al., (2009) found that activity in the ACCs coded for an interaction between the anticipated levels of effort and the levels of reward. Thus, there is evidence to support the claim that the overarching function of the ACCs may be to signal the net value of the rewards. This contrasts with the results of this study where activity in the ACCs was not modulated by the level of reward, only by the level of effort. The results of this study and that of Croxson et al. (2009) can be reconciled if one considers differences in the nature of the effort and the rewards in the two studies. Specifically, the reward values used in Croxson et al. (2009) were higher (25p and 5p) than those used in this chapter (16p and 4p). There was therefore a greater discrepancy between the HR and LR in that study than there was in this chapter. This greater discrepancy may have resulted in greater power to detect an interaction between reward and effort. As such, in this chapter, the ACCs activity may not have been modulated by reward, due to the low overall level of reward available in this study, which did not allow interactions to be identified.

5.5.2 Third-person effort-based valuation

To the best of my knowledge this study is the first to examine the neural basis of processing the amount of effort that another will expend to receive a reward. Key to the design of this study was that the processing of third-person reward and effort was not confounded with the subjects own reward and effort processing at the time of the instruction cues. This study was therefore able to examine whether activity in the ACCg varied with the anticipated net value of actions that were performed, and the reward received, by a third-person. Thus, the results of this study show that the ACC plays an important role in understanding the value of other's actions.

Broadly speaking the results of this study support claims made in chapter one, and in the previous two chapters, about the nature of information processing in the ACCg. The first claim was that the ACCs and the ACCg would process the discounted value of first-person and third-person actions respectively. The results of this study do not strictly support this interpretation; activity in the ACCg varied with third-person net-value, whereas activity in the ACCs was not attenuated by reward. Rather, activity in the ACCs varied with the amount of effort to be expended. Thus, the same information was not processed in these two areas. However, the ACCs was still engaged by effort-related information and, as suggested in section 5.5.1, this supports the notion that the area may still be important for the process of discounting rewards based on the level of effort to be expended. Thus, in general the results support the notion that both the ACCs and the ACCg are engaged in processing effort-related information for first-person and third-person actions respectively.

Anatomical evidence supports the notion that the ACCg may be engaged by effort-related information. The ACCg is strongly connected to area 24c' in the ACCs (Pandya et al., 1981; Vogt and Pandya, 1987), in which activity in this study varied with the level of effort. In addition, projections from the ACCg to the core of the nucleus accumbens partially overlap with projections from the ACCs (Devinsky et al., 1995; Haber et al., 1995). As stated above, both the ACCs and the nucleus accumbens are engaged when processing the value of rewards is discounted by an associated cost. The connections to these areas suggest that the ACCg has access to the effort-discounted reward values that are processed in the loop between the ventral striatum, ventral pallidum and the ACCs. Given that the ACCg also has connections to areas of the brain that are engaged in mentalizing (Vogt and Pandya, 1987; Frith and Frith,

2003), the connectional anatomy of the ACCg highlights this area as a region which may process other's predictions about the effort-discounted value of their actions.

Another interesting finding in this chapter was that the ACCg responded not only to a third-person's effort discounted reward values, but also to the undiscounted magnitude of a third-person's rewards. Whilst the aim of this study was not to test whether the ACCg processes the undiscounted magnitude of a third-persons rewards, these results support the notion that the ACCg processes similar information to the ACCs. There is considerable single-unit recording (Shidara and Richmond, 2002; Amiez et al., 2005; Sallet et al., 2007; Kennerley et al., 2009; Kennerley and Wallis, 2009b) and neuroimaging (Rogers et al., 2004; Knutson and Cooper, 2005; Hampton and O'Doherty, 2007) evidence to suggest that the ACCs codes for predictions about the magnitude of a reward, undiscounted by any other variable. In addition, the ACCg has strong connections to portions of the OFC (Carmichael and Price, 1995), PCC (Pandya et al., 1981) and intraparietal areas (Vogt and Pandya, 1987) that process rewarding stimuli. Thus, the functional properties of the ACCg highlight it as a candidate region for processing the predicted value of rewards that will be received by another.

Whilst this study tested specific hypotheses related to the ACC, results were also identified in other areas. Indeed, activity in several other areas either varied with the level of effort to be expended by another or responded differentially to high and low rewards that were to be received by another. Interestingly, there was very little overlap in the areas that processed these variables on the third-persons trials and those that processed them on the first-person trials. This result therefore suggests that different areas are engaged when processing other's predictions about the outcomes of effortful actions than those that are engaged when processing predictions about one's own effortful actions. However, two areas that have been implicated in processing first-person reward values were found to process third-person undiscounted rewards in this study. A cluster extending over the VmPFC and a separate region in the frontal pole showed a differential response to HR and low LR exclusively on the third-persons trials. As stated in the previous chapter, the VmPFC has strong connections to areas of the brain which are known to process the rewarding value of stimuli, such as the OFC, ventral striatum and the PCC and also to the ACCg (Pandya et al., 1981; Yeterian and Pandya, 1991; Morecraft et al., 1992). The VmPFC appears to respond to reward values flexibly, in many different task contexts. Neuroimaging research has shown that activity in this area varies with the predicted outcome of a decision, when the predicted value is based at least partially upon another's valuation (Behrens et al., 2008; Hampton et al., 2008). In the previous chapter, it was also shown that activity in this region coded for the predicted value that another was assigning

to their choice. These results therefore converge on the notion the VmPFC is engaged in processing reward predictions about upcoming responses. The results in this study extend upon previous findings, highlighting the VmPFC as important for processing another's prediction about the magnitude of reward they will receive for performing a series of actions.

5.5.3 Caveats and Limitations

Whilst the study provides new insight into how the value of other's effortful actions are processed, there are limitations to the design.

Firstly, the aim of this study was to examine how reward values are discounted by the amount of effort that is to be expended for their receipt. However, in the literature there is a debate as to what entails effort. In this study the effort task consisted of a series of targets cancelled by making finger movements to press buttons on a keypad. The preparation and execution of a button press therefore entailed physical effort. In addition, the number of cognitive operations that would be required to perform a large number of motor responses would be greater than when a small number of responses are required. Thus, the predicted amount of physical and cognitive effort expended would be confounded at the time of the instruction cues. However, recently, Kurniawaran et al. (2010) performed an fMRI experiment using a paradigm in which only physical exertion was modulated. They found that a portion of the striatum, in the putamen, responded to the level of physical effort. However, the reported activation was lateral in relation to the stratial activation reported in this chapter and in Croxson et al. (2009). Another study by Botvinick and colleagues (Botvinick et al., 2009) used a task where only 'mental effort' was manipulated and not physical effort. They showed that a portion of the ventral striatum was activated that was in close proximity to that activated in Croxson et al., (2009) and in this study. Thus, in this chapter, the modulatory effects of effort on the striatal signal may not be a result of the anticipated level of physical exertion. However, in this study, it is not possible to make claims, or speak to the debate, on what element of effort is processed in the ACC.

The second caveat relates to the length of time that subjects had to perform effortful actions. Tasks which examine the processing of different levels of effort, can sometimes confound the length of time that each series of effortful actions takes to perform with the amount of effort to be expended (Phillips et al., 2007; Wanat et al., 2010). In such cases, the reward may be

discounted by the length of the delay before receipt, rather than the level of effortful exertion (Day et al., 2010; Wanat et al., 2010). Alternatively, in tasks where the length of time for the effortful exertion is fixed, the task may become more difficult when the number of actions to be performed is greater. As such, the reward may be discounted by the task difficulty, but not by effort. To circumvent these problems in this study, subjects performed and observed another performing an effort task where the time period was fixed, but where subjects did not find the more effortful actions more difficult to perform. The accuracy of subjects in each of the four levels of effort was therefore an index of task difficulty. Subjects did not find the 12 cancellations more difficult than the two cancellations, indicating that performing 12 cancellations was not more difficult than 2 cancellations, but was more effortful.

5.5.4 Summary

In this chapter I tested two hypotheses relating to the valuation of rewards discounted by effort at the time of instruction cues. In support of the first hypothesis, activity in the ACCg varied with the net value of a third-persons reward discounted by effort. However, in contrast activity in the ACCs did not signal the net value of first-person rewards as hypothesised, rather it signalled the level of effort to be expended regardless of who was to perform the actions. The results of this study suggest that separate neural systems are used to value the effortful actions of oneself and others.

Chapter 6: Neural correlates of delayed rewards discounted by a social norm

6.1 Abstract

Understanding the behaviour of others requires the ability to understand that their subjective valuation of a reward may be distinct from your own. In some instances, one must use others' valuations to guide one's own behaviour, in order to conform to social norms. The previous chapter showed that the ACCg codes for the value of rewards that will be received by others, discounted by the effort expended for its receipt. Reward values are also discounted by the temporal delay before their receipt. This leads to the question: does the ACCg code for the value of rewards that are temporally discounted in a manner that conforms to a social norm? In this chapter, subjects performed delay-discounting trials, where they made a choice between a small reward, received immediately (£3), or a larger reward (£3.10 - £20) received following a delay (1 – 180 days). They were required to indicate either their own preferences or the preferences as dictated by a social norm on separate trials. Hyperbolic and exponential preference functions were fitted separately to the behaviour of subjects on the subjective trials (where they indicated their own preferences) and the normative trials (where they indicated the norm preference). The resulting idiosyncratic subjectively and normatively discounted reward values were fitted to the points in time when the subjects saw the delayed option and also to the points in time when they were instructed to indicate their choice. This design enabled two hypotheses to be tested. Firstly, that activity in the ACCs would vary with the subjectively discounted reward values and secondly that activity in the ACCg would vary with the normatively discounted reward values.

6.2 Introduction

In many situations individuals publically adapt their behaviour in order to conform to the attitudes and beliefs of the majority of others (Fehr and Fischbacher, 2004). Such conformity is often only public and individuals will continue to perform behaviours based on their own preferences when in private (Cialdini and Goldstein, 2004). There is considerable evidence that social norm conformity can occur despite considerable conflict between personal preferences and the preferences of the majority of others (Bond and Smith, 1996). This was illustrated by the seminal work of Asch (Asch, 1956). Asch (1956) showed that individuals conformed to a group answer on a perceptual judgement task, even when the norm judgement was a clear violation of both the subjects' own perception and the objectively correct answer. However, despite the fact that social norms can significantly bias an individual's decision-making, few studies have examined the processes that underpin normative economic decision-making.

Since the eighteenth century, much economic theory has been based around the concept that the objective value of a financial reward does not equate to the desirability of that reward to a subject (Bernoulli, 1954). The relationships between objective values and their subjective desirability can often be mathematically described using 'preference functions' (Mazur, 2001). One instance in which the objective value of a reward does not equate to the magnitude of the reinforcer subjective valuation was examined in the previous chapter, i.e. effort. Another variable that can discount the value of rewards subjectively is the delay before receipt (Mazur, 1985; Mazur et al., 1985; Green et al., 1994a). An individual who possesses a cheque for £10 which can be cashed in one month, might trade this for an £8 cheque they can cash immediately. The £10 has therefore been discounted such that its subjective value is now less than £8. This phenomena is known as temporal (or delay) discounting (Ainslie, 1974). Subjects can differ significantly on their preferences, with some individuals waiting only a few days for the £10, whereas others might wait many months. However, their behaviour can normally be explained by simple preference functions that contain a discount factor (Mazur, 1997; Richards et al., 1997). Discount factors relate the objective value of the rewards to the delay, idiosyncratically for each subject. Thus, the value of a subjectively discounted, delayed reward, motivates choices between differing financial options.

Rewards are discounted by delays before their receipt in many different species, for many different types of reinforcer. There is considerable evidence that humans, monkeys, rats, and many birds discount the value of primary reinforcers when their receipt is delayed (Ainslie,

1974; Ainslie, 1975; Freeman et al., 2009; Hwang et al., 2009; Jimura et al., 2009; Minamimoto et al., 2009; Dalley et al., 2011; Mar et al., 2011). Humans have also been shown to discount secondary reinforcers, such as delayed financial rewards (Rachlin et al., 1991; Green et al., 1994a; Green et al., 1994b; Kirby and Marakovic, 1996). Intriguingly, across species and across the different types of reinforcement, choices between a large delayed rewards and a smaller immediate reward can be explained by a hyperbolic preference function (Rachlin et al., 1991; Green et al., 1994a; Green et al., 1997; Kirby, 1997; Mazur, 1997; Richards et al., 1997; Mazur, 2001; Johnson and Bickel, 2002). However, to the best of my knowledge, no previous study has examined the behaviour of subjects on an economic delay-discounting task, when their choices must conform to a social norm.

In 2007, an influential paper was published by Kable and Glimcher (2007). They used fMRI to examine activity in the brains of subjects whilst they performed delay-discounting trials. They reported that the subjective preferences between a fixed immediate reward ($20) and a larger delayed reward (ranging between $20.25 and $110, delayed between 6 hours to 180 days) could be explained with a hyperbolic preference function. They found that activity varied parametrically with the idiosyncratically discounted subjective value of the delayed rewards in the ACCs, the ventral striatum, the posterior cingulate cortex (PCC), the intraparietal sulcus and superior portions of the OFC on the medial wall (Kable and Glimcher, 2007). Importantly, in the context of this thesis, these results highlight the ACCs as a candidate for processing the value of rewards discounted by one's own discount factor. Other neuroimaging studies have since shown support for the notion that the ACCs is engaged when discounting the value of delayed rewards (Peters and Buchel, 2009, 2010)

One of the aims of this thesis is to examine whether the ACCs and the ACCg process information that guides one's own decision-making and aids the understanding of other's decisions respectively. In the previous chapter, I examined activity in the ACC when rewards were discounted by the amount of effort required for their receipt. The study showed that the ACCs was engaged when subjects discounted the value of rewards they would receive, based on how much effort they would have to expend. The ACCg coded for the effort-discounted value of rewards, but did so when the effort that was to be expended, and the reward was to be received, by another. The previous chapter therefore showed that the ACCg codes for the net value of others' actions. Does the ACCg also process the value of delayed rewards when they are discounted by another's discount factor? This study examines whether activity in the ACCg codes for the discounted value of delayed rewards, when reward values are discounted in a manner that conforms to a (fictitious) social norm discount factor.

Subjects performed a delayed-response, temporal discounting task. On each trial subjects were presented with a choice between receiving £3 as payment for participation as soon as the experiment was complete or a larger reward (up to £20), the receipt of which would be delayed (between 1 day and 180 days). However, on 50% of trials subjects were required to indicate the social norm preferences, which they had learnt during a training session, instead of their personal preferences. Hyperbolic and exponential preference functions were fitted separately to the behaviour of the subject on both the subjective trials (where they indicated their own preferences) and on normative trials (where they performed the social norm behaviour). Each trial consisted of the presentation of the financial option (option cue) and then a trigger cue where they were required to indicate their choice between the delayed or immediate option. Activity was examined at both of these points in time on both the normative and subjective trials. Two hypotheses were tested: firstly, that activity in the ACCs would vary with subjective value of the discounted rewards and secondly, that activity in the ACCg would vary with the normative value of the delayed rewards.

6.3 Methods

6.3.1 Subjects

Subjects were sixteen healthy right-handed participants (aged between 18 and 30; 13 female), screened for neurological and psychological conditions. One (male) subject failed to respond on 30% of the trials and made multiple inter-slice head movements greater than 3mm. This subject was excluded from the analyses. Subjects were paid for their participation (see 'payment' below). The subjects were informed that a previous behavioural experiment had taken place with 102 participants. They were told that these participants received payment in the same manner that they had.

6.3.2 Experimental Design

The aim of this experiment was to examine the processing of delayed rewards, the value of which was either discounted subjectively, or in a manner that conformed to a social norm. Subjects performed a task over two consecutive days. On the first day, subjects learnt what they were told was the normative performance on a delay-discounting task. On the second day, subjects performed a delay-discounting task, where on separate trials they were required to indicate either their own "subjective" preferences or the "normative" preferences, which they had learned during training.

On each trial subjects were presented with a delayed option. Their task was to was to choose between this delayed option, or an immediate option of £3. The delayed options were always higher in magnitude than £3. The magnitude of the delayed options were £3.10p, £3.75p, £5, £8, £12 and £20 and were available at delays of 1 day, 15 days, 30 days, 60 days, 100 days and 180 days. Thus there were 36 different combinations of delay and magnitude that were used as the delayed options. The subjects were told that prior to the MRI experiment a behavioural experiment had been conducted on 102 subjects. Subjects were told that each of these 102 participants in the behavioural experiment had performed exactly the same delay-discounting task with the same delayed options as they would be presented with. They were told that, although preferences varied on the task, there was always at least 69% agreement (at least 70 participants) on which was the better option (to wait for the delayed reward or take the smaller reward immediately) for every delayed option they would see. During training,

subjects learnt this majority preference for each of the delayed options. This threshold of 69% ensured that subjects maintained the belief that they were learning a majority preference (or social norm). It should be noted that no such behavioural experiment was performed and the normative preferences that were learned were fictional. On the second day, during the scanning session, subjects performed trials where they chose between the delayed option or the immediate option. On some trials they were required to indicate their own preferences, on others they were told to choose what the social norm preference would be.

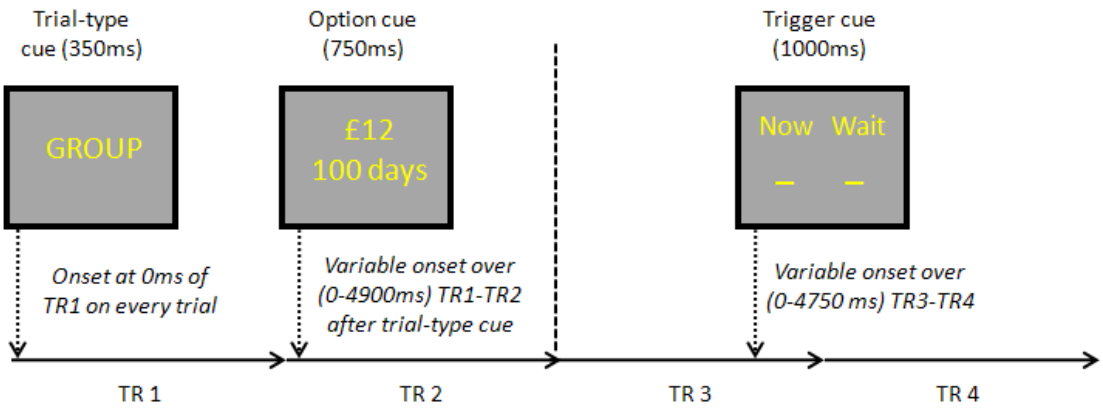

Fig. 6.1 Trial Structure. Each trial began with a trial-type cue that indicated whether the subjects should indicate their own preference or the normative preference. For a subjective preference the word "you" was the stimulus and for a normative preference the word "group" was the stimulus. Following a variable time period, the option cue was presented that indicated the delayed option. Subjects indicated their choice following the onset of a trigger cue. All stimuli were colour-coded, such that when the stimuli were yellow, the subject would indicate the norm preference and when the stimuli were white the subject would indicate their own preference.

6.3.3 Trial Structure

Subjects performed 216 delay-discounting trials (see fig.6.1) where they were presented with a delayed option and were required to choose between this option or the immediate option (£3). Each trial began with a trial-type cue (either the word "YOU" or "GROUP"), that indicated whether subjects were required to make subjective preferences ("YOU") or normative preferences ("GROUP") on the trial. Following the trial-type cue, after a variable delay,an option cue was presented that indicated the magnitude and delay of the delayed option. After a further variable delay, a trigger cue was presented where subjects were required to indicate their choice. A "now" stimulus was used to signify the £3 immediate option and a "wait"

stimulus was used for the delayed option. Subjects were required to indicate their choice at the time of the trigger cue, by pressing one of two buttons on a keypad. The trigger cue was presented for 1000ms, any responses before or after this time period resulted in the trials being classified as missed. Subjects were instructed to press the button which corresponded to "now" if their choice was the £3 immediate option or "wait" for the larger delayed option. In order to prevent subjects from preparing a specific digit response at the time of the option cue, the position of the "now" and "wait" stimuli were pseudorandomly organised, such that subjects could not predict which button would be "now" and which would be "wait" at the time of the option cue. All stimuli were also colour-coded to ensure that it was clear whether a subjective or normative preference was required on each trial. Yellow cues indicated that a normative choice should be made and white cues indicated that a subjective choice should be made. Subjects were presented with each of the delayed options six times. Thus, there were 108 subjective trials and 108 normative trials.

6.3.4 Procedure

6.3.4.1 Training

Subjects were trained in two phases one day prior to scanning. In the first phase the subject was seated in front of a monitor, with a response keypad. They were presented with a series of visual stimuli on the screen. They performed a series of delay-discounting trials where they were required to indicate their own preferences between a delayed or immediate option. Each trial consisted of an option cue (an amount of money and a delay period) and a trigger cue (two lines corresponding to two buttons on the keypad, with the words "wait" above one line and "now" above the other line). During this phase of training subjects performed 108 trials. This stage of the training enabled subjects to familiarise themselves with performing delay-discounting trials.

In the second phase of training, subjects performed a task where they learned the normative preferences for each delayed option. Each trial consisted of a delayed option cue, a trigger cue and a feedback cue. The feedback cue indicated the social norm preference for the delayed option on each trial. The cue was either the word "NOW" or "WAIT", which indicated whether the normative preference was for the immediate or delayed option respectively. Subjects were instructed to indicate the normative preferences on each trial. Thus, subjects were required to

monitor the feedback to improve their normative preference judgements on subsequent trials. The normative preferences learnt during this session were based around the behaviour of subjects during a pilot experiment (see "computational modelling" below for more details). During this session, subjects performed 108 trials, with the delayed options presented in the same order as they had been presented during the first phase of training. This session enabled subjects to learn the social norm preferences.

6.3.4.2 Scanning Session

On the day following the training session, subjects performed similar delay-discounting trials inside the MRI scanner. There were 216 pseudorandomly organised trials, 108 where they indicated their own preference (subjective trials) and 108 where they indicated the majority preference (normative trials). The same delayed reward options were used as during the training sessions. In this session, there was no feedback cue on the normative trials. Subjects were therefore required to recall the normative preferences they had learned during training. In this session, all stimuli were colour-coded, such that when stimuli were white, the subjects indicated their own preference and when stimuli were yellow they indicated the normative preference. There were 3 repetitions of each of the 36 delay-reward combinations for the subjective trials and the normative trials.

6.3.4.3 Payment

Subjects were told that their payment would be based on their choice on one of the trials during training or scanning. This payment would be determined by them selecting a number at random "out of a hat". The number would correspond to one of the trials during the training and scanning session and their payment would be based on whichever preference they indicated on that trial. This approach has previously been used to ensure that subjects are incentivised to accurately indicate their preferences on all trials (Kable and Glimcher, 2007; Pine et al., 2009). Subjects were informed that they would be paid by cheque for their participation. If on the selected trial they chose the delayed option, they would be paid that higher reward value, but the cheque would be dated such that it could not be cashed until the delay had passed. If they chose the immediate option then the cheque would be dated the same day. All of this information was provided to the subjects before the experiment, thus subjects were aware that their decisions during the experiment were real economic decisions.

The subjects were also told that their payment for the task would be based upon only their own preferences and not the normative preferences. It was possible that subjects could bias their choices on normative trials to make their payment for the experiment subjectively better. Thus, it was a necessity to pay subjects only based only on their subjective preferences in order to control the experimental design. Despite the subjects being told before the experiment that they would be paid based on a trial during either the scanning or training sessions, payment was based only on choices during the first phase of the training session. It was possible that subjects might alter their subjective preferences during scanning in order to conform or dissent from the normative preferences that they learned in the second phase of training. As the social norm in this study was fictional, it would have been unethical for the payment of a subject to have been biased by a normative behaviour created by the experimenter.

6.3.5 Computational Modelling

6.3.5.1 Behavioural Modelling

Previous research has shown that behaviour in delay-discounting tasks can be modelled using a number of different functions (Mazur, 1985, 1986; Green et al., 1994a; Green et al., 1997) that contain discount factors (free parameters that explain how rewards are idiosyncratically discounted by time). In this chapter, two models were compared separately in terms of their fit to the subjective preferences of the subjects and also the subject's behaviour on the normative trials. The first was a hyperbolic model (Mazur, 1986, 1997), in which the subjective value of a reward (V) is a function of its magnitude (M) and the delay (d):

(1)

In this model k is the discount factor, an idiosyncratic free parameter that discounts the magnitude (M) of the reward, such that the subjective value (V) is less than its objective magnitude. The value of k therefore reflects the extent to which a subject discounts a delayed reward, such that a high k decreases the value of the reward quickly as the delay becomes greater. There is considerable evidence that this model can explain preferences in delay-

discounting tasks (Kirby, 1997; Richards et al., 1997; Mazur, 2001; Johnson and Bickel, 2002; Madden et al., 2003; Kable and Glimcher, 2007).

As stated in chapter 4 (section 4.3), an important issue with studies using a computational modelling approach is whether the model reflects an accurate account of the computation's driving behaviour. To evaluate models thoroughly it is important to compare an alternative model that could also explain the behaviour on the task (Mars et al., in press). In this chapter, an alternative model was used that has previously been used extensively to examine delay-discounting behaviour (Kirby and Marakovic, 1996; Kirby, 1997; Johnson and Bickel, 2002; Madden et al., 2003). In this second model, the subjective value of the rewards (V) were discounted by an exponential function where:

(2)

In (2) the discounting effect of the delay is expressed as an exponential transform ( ) of the discount factor (k) multiplied by the delay period (d). As such, the magnitude of a reward (M) is idiosyncratically, but exponentially discounted by the length of delay before its receipt. The hyperbolic and exponential models were fitted separately to the preferences on the subjective trials and the choices on the normative trials. Thus, separate discount factors (k) could be estimated for the subjective and normative preferences.

6.3.5.2 Model estimation

To fit the two models to the behaviour of the subjects on both the subjective and normative trials, the softmax algorithm (Sutton and Barto, 1981; Sutton and Barto, 1998) was used. The softmax approach was employed separately for estimation of the normative and subjective discount factors. This method assigns a probability to the choices made by the subjects:

(3)

This equation converts the subjective values of the choices made by the subjects ( ) into a probability (, as a function of the value of both options. V_{o2} is the value of the delayed option and the value of the immediate option was always equal to 3 (£3 was always the immediate option). The coefficient β represents the stochasticity (or temperature) of the behaviour (i.e. the sensitivity to the value of each option). This algorithm therefore compared the value of the chosen option to the other options, the output is the probability of that option being chosen, given the value of the free parameters (k and β). The values were taken from the two models (see equation (1) and (2)) outlined above and fitted separately to both the subject's own preferences and also the behaviour on the normative trials. This allowed for comparisons to be made between the fit of the exponential and hyperbolic models for both the subjective and normative behaviours.

Within the two preference functions, and also in the softmax algorithm, there are free parameters which need to be estimated. To identify the optimal set of free parameters for each preference function, the values of the discount factors (k) were varied from 0.00001 to 0.2 in steps of 0.00001 and β was varied from 0.01 to 2 in steps of 0.01. The output of the softmax algorithm is a series of probabilities, based on the values of each of these parameters and the choices made by the subject. These parameters were varied separately both for the subjective and normative preference behaviours. By varying the parameters, the probabilities output by the softmax algorithm differ. To select the parameters that best fitted the choice behaviour, given the preference function, a maximum likelihood approach was used. By using a maximum likelihood algorithm it was possible to maximise the probabilities of the choices

made by the subjects and esimate the values of each of the parameters that produced the behaviour. The likelihood of the chosen responses were found using:

(4) (n)

Where the likelihood of each set of parameters (L) is determined by the log of the probability of the chosen response ( ) at trial *n,* according to the model. If the model perfectly predicts the actions, the probability of every chosen response would equal 1 and L would be 0. As the probabilities become less than 1 the log-likelihood L assumes negative values. The best fitting parameters were then selected using:

(5)

This identified the set of parameters for which L was closest to 0, i.e. the best fitting parameter set. Where is the parameter set and L is the log-likelihood. In this study, it was also important to determine which of the two preference functions fit better to the subjective preferences and also the choices on the normative trials. To determine which was the best fitting model, the log evidence for the best fitting parameters in each subjects were entered into a 2x2 repeated measures ANOVA. One factor was the nature of the preference (normative or subjective) and the second was the preference function (hyperbolic or exponential). It is important to note that more sophisticated methods of model comparison exist, such as the Akaike information criterion or the Bayesian information criterion (Mars et al., in press). However, these approaches are only appropriate when comparing models with different numbers of free parameters. Here, both the hyperbolic and exponential functions required the estimation of the same number of free parameters. As such, comparing the models based on their log-likelihood is a suitable approach.

During the training session subjects learnt the normative preferences through feedback. The normative behaviour that was presented as the feedback cues was actually the predicted choices of a hyperbolic model, with a fixed discount factor of 0.02173 and a temperature (β) of 1. These values reflected a mean of 8 subjects' behaviour on the delay-discounting task during a pilot experiment. The hyperbolic model was chosen for this purpose, rather than the exponential model, as previous studies have shown that the hyperbolic model accounts for

discounting behaviour better than an exponential model (Green et al., 1997; Kirby, 1997; Mazur, 2001; Madden et al., 2003; Kable and Glimcher, 2007).

6.3.6 Experimental timing

An important feature of the study was that activity was time-locked to two different time points on each trial. In order to do this, a variable delay was introduced between the delayed option and the trigger cue. This allowed for BOLD activity time-locked to both the option cues and the trigger cues to be isolated, without the contaminating effects of either prior or subsequent trial events (Ramnani and Miall, 2003, 2004). Events in each trial took place across four TRs (0s–12s; TR=3s). To optimally sample the cues of interest, randomly varying intervals between the scan onset and these cues were introduced over the first and second TRs for the option cues and over the third and fourth TRs for the trigger cues. This achieved an effective temporal sampling resolution much finer than one TR for the conditions of interest. The intervals for both the trigger cue and option cue were evenly distributed, ensuring that EHRs time-locked to each cue was sampled evenly across the time period following an event. Separate jitters were employed for the option cues and the trigger cues on the subjective trials and on the normative trials.

6.3.7 Functional imaging and analysis

6.3.7.1 Data acquisition

878 EPI scans were acquired from each participant. 38 slices (10% distance factor) were acquired in an ascending manner, at an oblique angle (≈30˚) to the AC-PC line to decrease the impact of susceptibility artefact in the subgenual ACC (Deichmann et al., 2003). A voxel size of 3×3×3 mm (20% slice gap, 0.6 mm) was used; TR=3s, TE=32ms, flip angle=85°. The functional sequence lasted 51 minutes. High resolution T1-weighted structural images were also acquired at a resolution of 1×1×1 mm using an MPRAGE sequence. Immediately following the functional sequence, phase and magnitude maps were collected using a GRE field map sequence (TE_1 = 5.19ms, TE_2 = 7.65ms).

6.3.7.2 Image preprocessing

Scans were pre-processed using SPM8 (www.fil.ion.ucl.ac.uk/spm). The EPI images from each subject were corrected for distortions caused by susceptibility-induced field inhomogeneities using the FieldMap toolbox (Andersson et al., 2001). This approach corrects for both static distortions and changes in these distortions attributable to head motion (Hutton et al., 2002). The static distortions were calculated using the phase and magnitude maps acquired after the EPI sequence. The EPI images were then realigned, and coregistered to the subject's own anatomical image. The structural image was processed using a unified segmentation procedure combining segmentation, bias correction, and spatial normalization to the MNI template (Ashburner and Friston, 2005); the same normalization parameters were then used to normalize the EPI images. Lastly, a Gaussian kernel of 8 mm FWHM was applied to spatially smooth the images in order to conform to the assumptions of the GLM implemented in SPM8.

6.3.7.3 Event definition and modelling

In this study, two GLM analyses were performed to investigate activity that varied parametrically with the subjective and normative values of temporally discounted rewards. The first GLM was employed to examine activity that varied with the normatively and subjectively discounted values. Regressors which modelled the subjective and normative values were orthogonalised with respect to regressors that modelled the delay periods. Thus,

activity that varied with the value parameters was not shared with the effort parameters. In the second GLM the delay parameters were orthogonalised with respect to the subjective and normative values parameters.

Each GLM design matrix contained regressors modelling:

- Trial-type cue (informing the subject whether they should indicate their own preference or the majority preference on the trial)
- Subject option cue (the delayed reward option on the trials where the subjects indicated their own preferences)
- Norm option cue (cuing the subject to indicate the normative preference)
- Subject trigger cue (the trigger cue when the subject indicated their own preference)
- Norm trigger cue (cuing the subject to indicate the normative preference)
- Missed trials (a regressor modelling the onsets of the option and trigger cues of missed trials)

Regressors were constructed for each of these events by convolving the event timings with the canonical Haemodynamic Response Function (HRF). The residual effects of head motion were modelled in the analysis by including the six parameters of head motion acquired during preprocessing as covariates of no interest. In addition to the regressors defined for the event types outlined above, each GLM also contained regressors which were first order parametric modulations of the option cue and trigger cue events. These modulators scaled the amplitude of the HRF in line with either the length of the delay or the discounted reward values. Separate parametric modulators were created for the delay and discounted reward values on the subjective trials and the normative trials. It was not possible to use the undiscounted reward values as parametric modulators, due to the colinearity with the subjective (or normative) reward values (i.e. the subjective reward values of subjects with low discount factors were highly correlated with the undiscounted reward values). To examine activity that varied with the subjective and normative values, the regressors which constitute the parametric modulators of the two values were orthogonalised with respect to the delay parametric modulators. As such, activity varying with the subjective reward parameters at the time of both the subject trigger cue and the subject option cue events could not be explained by activity related to the delay. Likewise, activity that varied with the value of normative reward parameters at the time of both the norm trigger cue and the norm option cue events could not be explained by delay related activity. This approach allowed for two hypotheses to be tested:

(i) that activity in the ACCs would vary with the subjectively discounted reward values on the subjective trials, and (ii) that activity in the ACCg would vary with the normatively discounted values on the normative trials.

A second GLM was also constructed which contained the same regressors as outlined above. However, in this GLM the delay parameters were orthogonalised with respect to the subjective and normative reward parameters. This GLM therefore allowed for activity which varied with the length of the delay to be examined.

6.3.7.4 Group analysis, Contrasts and Thresholding

Random effects analyses (Full-Factorial ANOVA) were applied to determine voxels significantly different at the group level. SPM{t} images from all subjects at the first-level were input into second-level full factorial design matrices. F-contrasts were conducted on the regressors for the delay and reward parameters in each of the GLMs. These contrasts identified voxels in which activity varied parametrically in the manner predicted by the subjective or normative value or the delay parameters. Separate corrections for multiple comparisons were used for the ACCg, ACCs and the whole brain. To examine activity across the whole brain, FDR correction ($P<0.05$) was applied. In contrast, activity in the ACCg and ACCs was corrected for by using an 80% probability mask of each region (see chapter two for a description). In addition, due to the close similarity of the design to that of Kable and Glimcher (2007), small volume corrections were applied around the peak coordinates from their analysis (see appendix E for table of coordinates). 8mm diameter spheres were created around those coordinates, and then used as one mask containing the spheres from all regions.

6.4 Results

6.4.1 Behavioural results

Subjects performed a delay-discounting task, in which they indicated either their own preferences or preferences dictated by a social norm, on separate trials. Hyperbolic and exponential preference functions were fitted separately to the subjects' behaviour on the subjective trials and the normative trials. Log-evidence was calculated for the best-fitting set of parameters for each preference function for both the normative and subjective behaviours. A repeated measures 2x2 ANOVA was conducted examining log-evidence for the two preference functions (Hyperbolic or Exponential) and two types of trial (Subjective or Normative) across subjects. No interaction between the log evidence for the different levels of trial type and preference function was found ($F(1,14) = 3.874$, $p = 0.069$). However, there was a significant main effect of preference function ($F(1,14) = 17.554$, $P<0.001$), with no main effect of trial type ($F(1,14) = 4.391$, P = 0.055). Examination of the log-evidence (see fig.6.2) shows that the hyperbolic model had a significantly lower log-evidence than the exponential model for both the subjective and normative trial types. Thus, the hyperbolic model explained preference behaviour in both the subjective and normative conditions better than the exponential model. In addition, there was no difference in log-evidence between the hyperbolic model on the normative and subjective trials. As the hyperbolic model was a better fit of the data in both conditions, this model was analysed further behaviourally and the values it output were used as the values for the parametric modulators in the fMRI analysis.

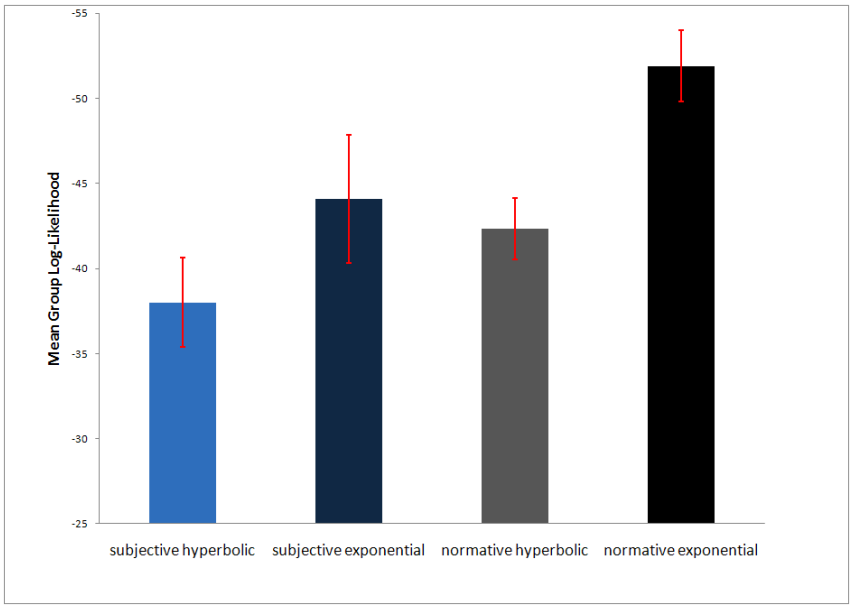

Fig 6.2 Mean Log-evidence across subjects for the hyperbolic and exponential preference functions on both the subjective and normative trials. Log evidence values closer to zero indicate a better fit to the data.

This study assumed that subjects acquired a model of a social norm discount factor (k) during training, in order that they could reproduce the normative behaviour during scanning. It was therefore pertinent to examine whether the estimated discount factor (k) of the behaviour of subjects in the normative condition, was not significantly different from the discount factor which underpinned the social norm preferences that they learned during training. A paired samples *t*-test was conducted to examine whether the best fitting values of k across subjects (mean k= 0.029) was different from the k which was used to develop the normative preferences (actual k = 0.02173). No significant difference was found between the mean k of the subject's behaviour on the normative trials from the k which underpinned the norm preferences ($t(14) = 0.919$, $p = 0.374$). This indicates that the model the subjects' formed of the normative discount factor was not significantly different from the discount factor that underpinned the social norm. Another pertinent question, both behaviourally and for the interpretation of the fMRI results, is whether subjects' results could be explained by their divergence from the social norm i.e., does the divergence of the normative discount factors from the social norm discount factor (k = 0.02173) correlate with the divergence of the subjective discount factors from the social norm. A Spearman's rank correlation revealed no significant correlation between the divergence of the normative and subjective discount factors from the actual social norm discount factor ($R^2 = 21.9$, $p = 0.078$, one tailed). Thus, subjects' performance of the social norm preferences on the normative trials was not correlated with their own subjective preferences or divergence from the social norm.

In addition to examining the discount factor, it was also pertinent to examine how accurate the subjects were at predicting the choices that would have been made based on the social norm discount factor. All subjects were better than chance at predicting the normative preference (i.e. the number of trials where the chosen option was the same as the actual normative preference; mean across subjects 82.59%, S.D ±7.11). A paired samples *t*-test showed that the accuracy of subjects in terms of performing the normative choice preferences was significantly different ($t(14) = 17.76$, $P<0.001$) from chance (50%). This suggests that subjects' discounted reward values on the normative trials hyperbolically, and in a manner that closely fitted with the actual normative preferences.

An additional consideration was whether the hyperbolic model more accurately predicted the choices of subjects in the normative or subjective conditions. To examine this issue a paired *t*-test was conducted on the accuracy of the model (percentage of the choices accurately predicted by the model) on the subjective trials (mean = 86.7%) compared to the normative trials (mean = 86.1%). No significant difference in model accuracy was identified between the

two types of trials ($t(14) = 0.368$, $p = 0.718$). In addition, as reported above, there was no difference in log-evidence between the fit of the hyperbolic model in the subjective and normative conditions. These results suggest that the hyperbolic model did not differ in accuracy between the two types of behaviour.

The behavioural results therefore indicate that, in line with previous studies, subjects personal preferences on a delay-discounting task could be explained with a hyperbolic preference function. In addition, the results indicate that subjects could learn and reproduce normative preferences on a delay-discounting task and this behaviour could also be explained with a hyperbolic function.

6.4.2 Imaging results

6.4.2.1 Subjective reward value

To examine activity that varied with the subjective value of rewards, and not the normative value, at the time the delayed option, idiosyncratic parametric modulators were fitted to the option cues of each subject. This parameter modelled the subjective value of a reward, based on the subject's own discount factor in the hyperbolic model. To examine whether activity in areas previously implicated in processing the subjective value of rewards varied with this parameter in this study, a small volume correction was applied around the peak coordinates of Kable and Glimcher (2007). Activity was found to vary with the subjective reward values in the three areas, the ventral striatum/nucleus accumbens (-14,10,-8, $F(1,126) = 25.21$, $Z = 4.64$, $p<0.05$FDR), in the ACCs (-4,46,14; $F(1,126) = 6.96$ $Z = 2.93$, $p<0.05$svc), on the borders of the rostral MCC/caudal ACC (at the borders of areas 24c, 24c', 32 and 32', see fig. 6.3) and the posterior Superior Temporal Sulcus (64, -54, 16; $F(1,126) = 12.69$, $Z = 3.24$, $p<0.05$svc). Correcting for multiple comparisons across the whole-brain did not identify activity in any areas that varied significantly with the subjective value of reward at the time of the option cue.

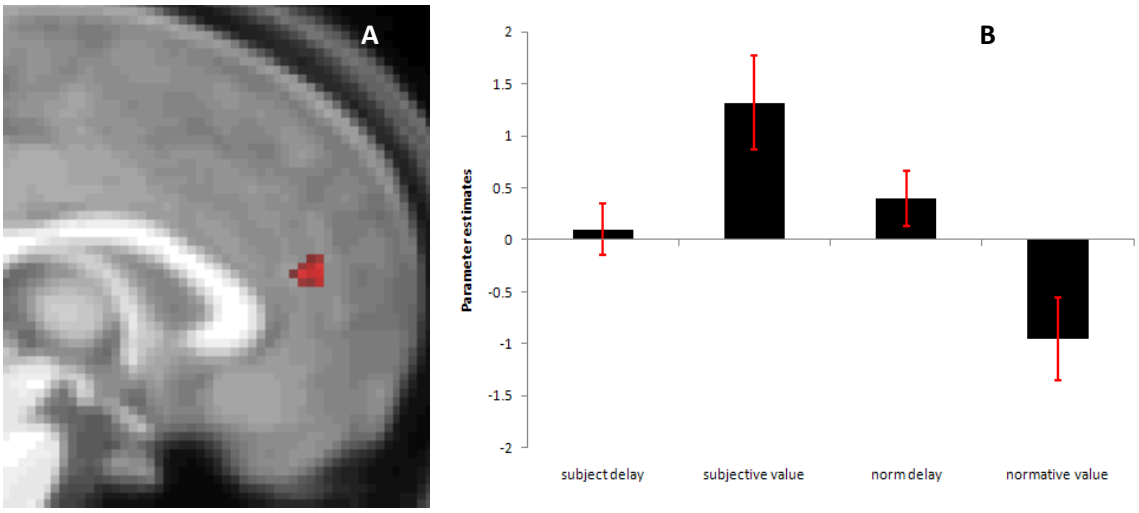

Fig 6.3 Subjective reward values at the time of the delayed option cue. (A) Activity shown in the ACCs (-4, 46, 14) that varied with the subjective valuation of the delayed option. (B) Beta coefficients of activity from the peak ACCs voxel. The parameter estimates are taken from the GLM in which the subjective and normative value parameters were orthogonal to the delay parameters.

In addition to examining activity at the time of the option cue, activity was also examined at the time of the trigger cue. A whole brain analysis revealed activity in the ventral striatum (-2, 6,-4; F(1,140) = 18.24, Z = 3.97, p<0.05 svc, see fig.6.4) and a cluster extending over subgenual cortex (12, 40, 0; F(1,140) = 22.06, Z = 4.37 p<0.05 svc) and the superorbital sulcus, often referred to as Ventromedial Prefrontal Cortex (VmPFC).

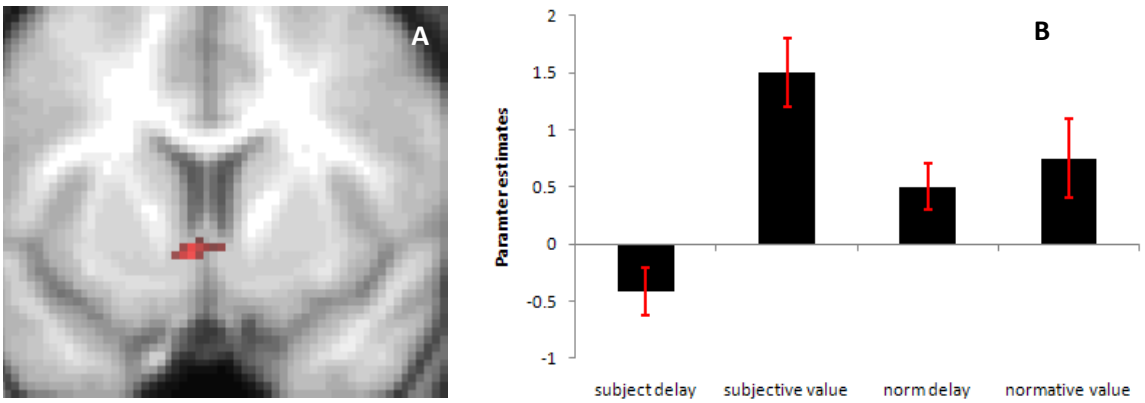

Fig 6.4 Subjective reward values at the time of the trigger cue. (A) Activity shown in the ventral striatum (-2, -6,-4) that varied with the subjective valuation of the delayed option. (B) Beta coefficients of activity from the peak striatal voxel. The parameter estimates are taken from the GLM in which the subjective and normative value parameters were orthogonal to the delay parameters.

6.4.2.2 Normative reward value

To examine activity that varied exclusively with the normative reward values, idiosyncratic parametric modulators were fitted to the option cues on the normative trials. The values of this parameter were taken from the output of the hyperbolic model on the normative trials. A whole-brain analysis revealed activity in several areas which varied with the normative reward values (see table 6.1). No voxels in the ACCg were found to vary with the magnitude of the normatively discounted reward at the time of option cue. However activity was found in the paracingulate cortex (0, 20, 46; F(1, 84) = 37.65, Z = 5.44, p<0.05FDR), a region heavily implicated in the processing of social information (See fig.6.5). In addition to activity at the time of the option cue, activity was also examined at the time of the trigger cues. No brain area was found when correcting for multiple comparisons across the whole brain. However, activity in the ACCg (putatively midcingulate area 24b'; see fig.6.5) survived small volume correction of the ACCg (4, 24, 28; F(1,140) = 2.95, Z = 2.90, p<0.05svc).

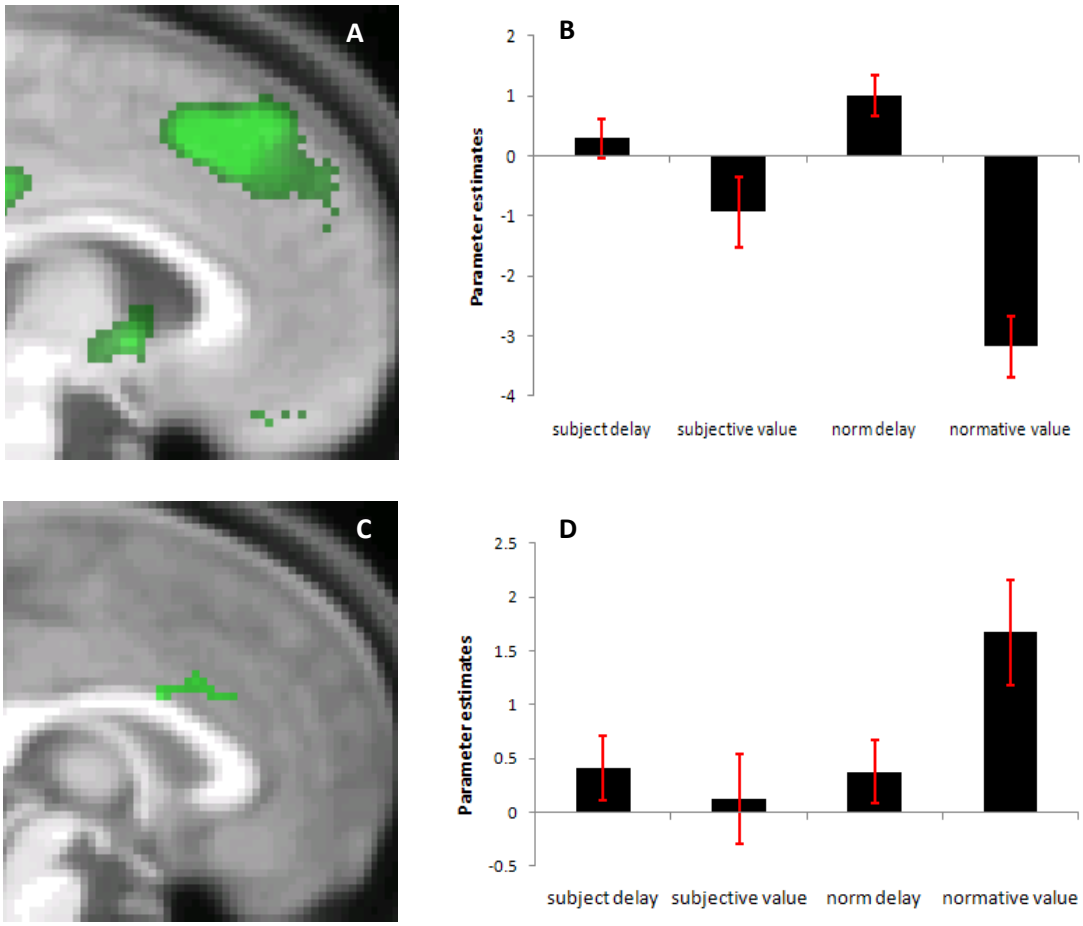

Fig 6.5 (A) Activity shown in paracingulate cortex that varied with the normative value of the delayed option at the time of the option cue. Activity also shown in the Subcallosal gyrus and the posterior cingulate gyrus/retrosplenial cortex. (B) Beta coefficients of activity from the peak paracingulate voxel. (C) Activity shown in the ACCg that varied with the normative values of the delayed option at the time of the trigger cue. (D) Beta coefficients of activity from the peak ACCg voxel. The parameter estimates are taken from the GLM in which the subjective and normative value parameters were orthogonal to the delay parameters.

Anatomical region	***MNI Coordinate***	***Z-value***
Parietal		
Intraparietal Sulcus	36, -60, 50	5.93
Short Insular Gyrus	-32, 22, 0	3.75
Frontal		
Paracingulate Sulcus	0, 20, 46	5.67
Middle frontal Gyrus	42, 32, 18	4.56
Orbital Gyrus	42, 51, -6	3.67
Limbic		
Posterior Cingulate Gyrus	-4, -38, 34	4.86
Hippocampus	-20, -34, -10	4.17
Striatum		
Ventral Caudate/Subcallosal Gyrus	4, 4, -4	4.05
Occipital		
Cuneus	-14, -80, -8	3.89

Table 6.1 Areas in which activity varied with the normative reward values at the time of the option cue. All regions survived whole brain FDR (P<0.05) correction for multiple comparisons. All contrasts were F(1,140) contrasts on the normative reward parametric modulators, exclusively masked by the subjective parametric modulator (p<0.05unc).

6.4.2.3 Subjective and normative reward values.

To identify whether activity in any region varied with the value of rewards discounted by delay, regardless of the whether they were subjectively or normatively valued, conjunctions were performed between the normative and subjective value parameters. At the time of the option cue, no area was identified in a whole-brain analysis. Small volume correction of the ACCs identified a region in the anterior portion of the ACC (putatively area 32'). Uncorrected this cluster extended across a large portion of the Ventromedial Prefrontal Cortex (-8, 44, 16; F(2, 126) = 6.58, Z = 2.38, p<0.05svc). Activity in this area varied with the value of discounted rewards, regardless of whether they were normatively or subjectively discounted. At the time of the trigger cue, a different portion of the ACCs (in posterior midcingulate area 24c') (-12, 14, 38; F(2,126) = 7.28, Z= 2.42, p<0.05svc) varied with both the subjectively and normatively discounted reward values.

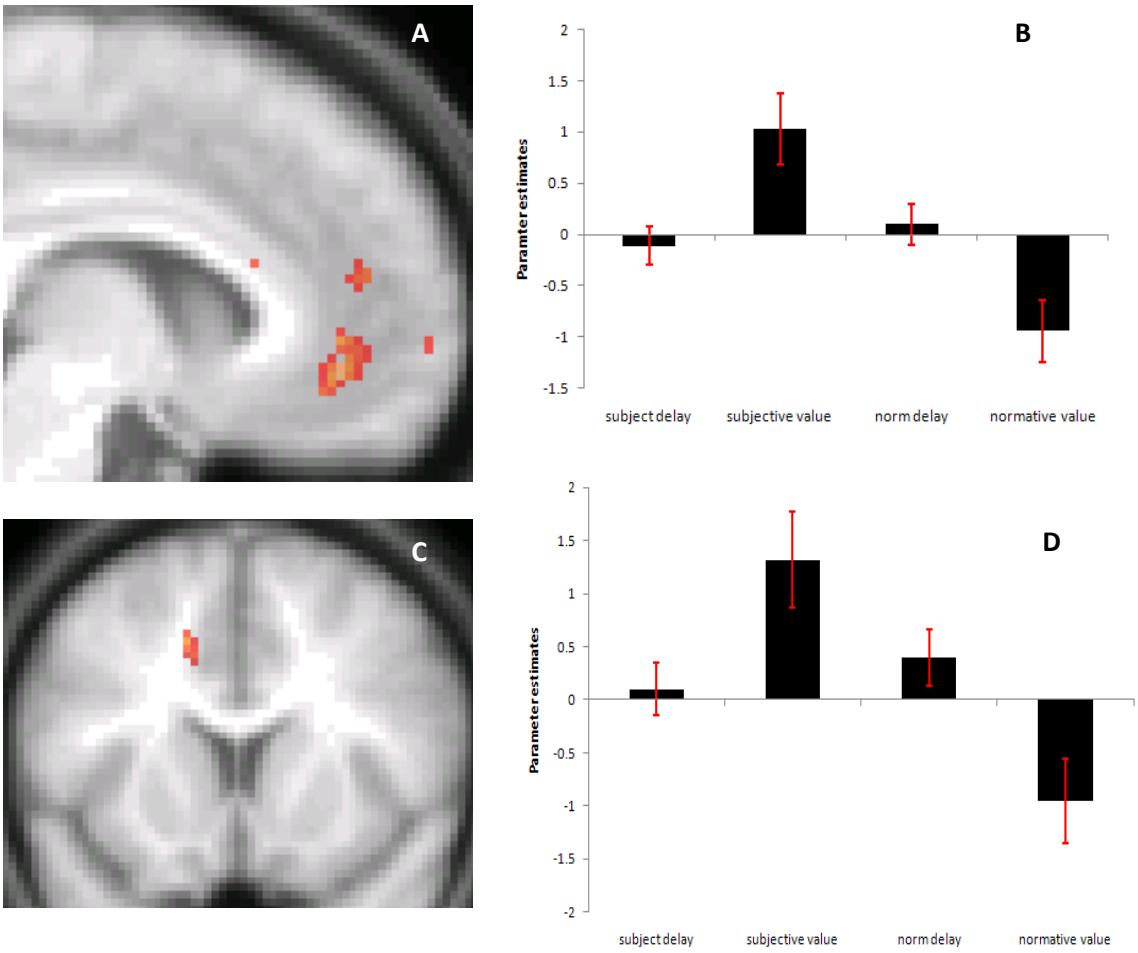

Fig 6.6 (A) Activity shown in Ventromedial Prefrontal cortex (-8, 44, 16 ; VmPFC) that varied with both the normative and subjective valuations at the time of the option cue. This image is shown uncorrected for display purposes (B) Beta coefficients of activity from the peak VmPFC voxel. (C) Activity shown in the ACCs (-12, 14, 38) that varied with the normative and subjective values of the delayed option at the time of the trigger cues. (D) Beta coefficients of activity from the peak ACCs voxel. The parameter estimates are taken from the GLM in which the subjective and normative value parameters were orthogonal to the delay parameters.

6.4.2.4 Normative and Subjective Delays

In addition to the subjective and normative reward values, the delay parameter was fitted to the option cue and the trigger cue events. No activity was found to vary with the length of the delay on the subjective trials at either the time of the option cue or the trigger cue. However, several areas were found to vary with length of delay on the normative trials at both the time of the option cue (see table 6.2 for results) and the time of the trigger cue (see table 6.3 for results).

Anatomical region	***MNI Coordinate***	***Z-value***
Parietal		
Superior Parietal Gyrus	-8, -68, 28	4.74
Sub parietal Sulcus	-14, -48, 34	3.90
Cingulate		
Posterior Cingulate Gyrus	6, -20, 34	3.42
Posterior Cingulate Sulcus	-2, -22, 56	3.88
Frontal		
Middle frontal Gyrus	-30, 18, 50	3.89

Table 6.2 Areas responding to the length of delay at the time of the option cue on the normative trials. Whole brain corrected for multiple comparisons (p<0.05 FDR). All contrasts were F (1,126) on the normative delay parametric modulator masked by the subjective delay parameter.

Anatomical region	***MNI Coordinate***	***Z-value***
Parietal		
Superior Parietal Gyrus	10, -80, 48	4.34
Cerebellum		
Crus I/II	34, -78, -26	4.21
Frontal		
Olfactory Sulcus	-20, 24, -12	4.40
Inferior Frontal Gyrus	-56, 24, 26	4.35
Gyrus Rectus	2, 50, -22	4.17
Superior Frontal Gyrus	-2, 60, 32	3.98
Middle Frontal Gyrus	48, 34, 28	3.82

Table 6.3 Areas responding to the length of delay of the delayed option at the time of the trigger cue. Whole brain corrected for multiple comparisons (p<0.05 FDR). All contrasts were F (1,126) on the normative delay parametric modulator masked by the subjective delay parameter

6.5 Discussion

This study tested two hypotheses related to the processing of rewarding values, discounted by the temporal delay before their receipt. The first hypothesis was that activity in the ACCs would vary with the value of rewards hyperbolically discounted by a subject's own discount factor. The second hypothesis was that activity in the ACCg would vary with the value of rewards hyperbolically discounted by a social norm discount factor. In accordance with both hypotheses, subjects behaviour was better explained by a hyperbolic preference function, than an exponential preference function, both when indicating their own preferences and normative preferences. In line with the first hypothesis, activity in the ACCs varied with the magnitude of the delayed reward at the time that the delayed option was presented. In line with the second hypothesis, activity in the ACCg varied with the magnitude of the normatively discounted delayed rewards, at the time of trigger cue. Activity was also found in the paracingulate cortex that varied with the normatively discounted rewards values, at the time the delayed option was represented. In addition, activity in two different portions of the ACCs varied with the magnitude of the discounted rewards regardless of whether it was subjectively or normatively discounted. These results support the claim that the ACC is important for processing the value of delayed reward. Activity in the ACCs signals the value of a reward regardless of the source of the temporal discounting, whereas activity in the ACCg codes for the valuations of others.

6.5.1 Temporally discounted value signals in the ACC

This study showed that activity in different portions of the ACC code for both subjectively and normatively discounted rewards, highlighting the ACC as a whole as an important area in processing discounted reward values. Interestingly, activity in the ACCs (in posterior MCC, putatively in the RCZ) and the ACCg occurred at the time that subjects were required to make their response and not at the time that the delayed option was first presented. This result would suggest that the ACCg, and also the ACCs, are important for the process of assigning value to actions, rather than discounting the value of a delayed option when it is first presented. The notion that the ACCg is engaged in processing values when they can be assigned to actions is supported by the work in previous chapters. In chapter five, activity in the ACCg signalled the value of another's rewards that were discounted by the number of

effortful actions they would be required to perform. Thus, in both this chapter and the previous, the ACCg coded the value of actions that were devalued by a cost. Chapter 4 showed that activity in the ACCg occurs when another's prediction of the value of their action is erroneous. Thus, whilst no study has stringently tested whether the role of the ACCg is to process the value that others place on actions, the evidence in this thesis supports this claim.

The results in this chapter extend beyond those of previous fMRI studies which have investigated how other's valuations of rewards guide first-person behaviour. Firstly, the results indicate that the ACCg does not just process the value of rewards discounted be the expenditure of effort, or the volatility of other's reward-related decision-making (Behrens et al., 2008), but also processes the normative valuation of temporally discounted rewards . Secondly, the results indicate that the ACCg also processes other's reward valuations (in this case a norm valuation) when they are guiding first-person behaviour. This study is not the first to show that the ACCg plays a role in guiding first-person behaviour. Behrens et al., (2008) showed that activity in the ACCg varied with the volatility of another's advice about the value of two rewarding options. Moreover, the level of activity in the ACCg correlated with the degree to which subjects' decision-making was biased by their advice. Subjects' whose activity in the ACCg was increased but decreased in the ACCs were more influenced by others' valuations that those with the opposite profile. However, the study reported here is distinct from Behrens et al. (2008) in two ways. Firstly, in this chapter the signal in the ACCg reflected a normative valuation of a reward, not just the valuation of one individual. Secondly, the activity in the ACCg varied with the values of a temporally discounted rewards and not the value of another's advice in terms of its volatility, as in Behrens et al. (2008). The results reported in this chapter and the study of Behrens et al. (2008) therefore support the notion that the ACCg codes for other's valuations, when a subject is required to use that valuation to guide their own behaviour.

Whilst activity in the RCZ occurred at the time of the trigger cue, activity in a more anterior region (at the border between MCC and ACC) occurred at the time the delayed option was presented. Interestingly, this activity varied both with the normatively and subjectively valued delayed rewards. Neurons in this region have been shown to be sensitive to the value of rewards (Amiez et al., 2005; Kennerley et al., 2009; Hayden and Platt, 2010). The cluster also overlapped with the activation reported in Kable and Glimcher (2007), who also showed that activity in this area varied with the subjective value of temporally discounted rewards. However, unlike in this previous studies, activity in this study could be examined at the point in time when the delayed option was presented and also separately at the time that an action

was triggered. Thus, this study was able to extend upon the results of Kable and Glimcher (2007), showing that the anterior MCC is processes temporally discounted rewards at the time a delayed option is valued, but activity in the RCZ processes activity when such signal motivate actions.

6.5.2 Social Norm Discounting

This study showed that areas typically implicated in processing the subjective value of delayed rewards also process normatively discounted delayed rewards. The Intraparietal sulcus (IPS), OFC, the middle frontal gyrus and the PCC varied parametrically with the value of normatively discounted delayed rewards at the time of the option cue. Each of these areas has previously been implicated in processing discounted reward values (Cardinal, 2006; Doya, 2008; Rushworth et al., 2009). Single-unit recording studies have shown neurons in the IPS (Louie and Glimcher, 2010), OFC (Roesch et al., 2006; Roesch et al., 2007) and the DLPFC (Kim et al., 2008) which are sensitive to reward values discounted by a temporal delay before their receipt. Neuroimaging studies also show activity in each of these areas and also in the PCC when subjects discount delayed rewards (Apicella et al., 1991; Knutson et al., 2000; McClure et al., 2004; Kable and Glimcher, 2007; Luhmann et al., 2008; Peters and Buchel, 2009; Pine et al., 2009; Kable and Glimcher, 2010; Prevost et al., 2010). This would suggest that discounting rewards based on a social norm discount factor, requires the engagement of a system that is also involved in processing first-person valuations of delayed rewards. There is an obvious caveat to this interpretation, in that these areas did not vary with the subjective reward values in this experiment. However, one possible explanation for this is that there is considerable variability in the type and magnitude of rewards and delay periods that have been used in the previous studies. Single-unit recording studies in non-human primates and lesion studies in rats use primary reinforcers, rather than the secondary reinforcers, such as the financial incentives used in this chapter (Richards et al., 1997; Denk et al., 2005; Roesch et al., 2006; Rudebeck et al., 2006b; Kobayashi and Schultz, 2008; Wanat et al., 2010). Neuroimaging studies investigating delay-discounting which have used financial incentives, have used different levels of reinforcement to those employed in this chapter (Kable and Glimcher, 2007; Pine et al., 2009; Prevost et al., 2010). Other studies have also used non-financial stimuli to act as reinforces (Prevost et al., 2010). In addition to variability in the incentives used, there is also considerable variability in the durations before which the reinforcer is received in both the

animal and human studies. Such variability in the type and range of reinforcement and the durations before the receipt of the reinforcer, may in part explain why this study failed to find activity on the subjective valuation trials, in the areas which have previously been implicated in processing first-person discounted rewards. It is therefore pertinent not to interpret the absence of such signals in this study as indicative of these areas not being engaged in subjective reward valuation. Thus, despite the caveat, it is still reasonable to suggest that normative delay-discounting may rely in part on systems which are engaged in subjectively valuing delayed rewards.

The results of this study suggest that the ACCg and the paracingulate cortex may be important for processing other's valuations of delayed rewards. Anatomical evidence suggests that the paracingulate cortex and the ACCg may exchange information with these areas which are engaged when valuing delayed rewards. Paracingulate cortex (at the border between area 8/9 and 32' on the medial surface) in the macaque brain, shows strong reciprocal connections with area 14 (medial OFC; Petrides and Pandya, 2006, 2007), area 24c' (the CMAr/RCZ; Vogt and Pandya, 1987; Petrides and Pandya, 2006, 2007), areas 23c'/31 (PCC; (Petrides and Pandya, 2006, 2007) and it also send projections to the ventral portions of the striatum (Apicella et al., 1991; Yeterian and Pandya, 1991; Kunishio and Haber, 1994; Haber et al., 1995). As outlined above, each of these areas has been implicated in processing temporally discounted reward values previously. Paracingulate cortex and the ACCg are also reciprocally connected (Vogt and Pandya, 1987; Yeterian and Pandya, 1991; Morecraft et al., 1992; Morecraft and Van Hoesen, 1998). Thus, these two areas are connected to each other and also to areas of the brain which process discounted reward values when the valuation is subjective. The connectivity profiles of these two regions therefore suggest that areas which are specialised for processing social information exchange information with areas which guide reward-related first-person decisions. This supports the view that behaviour guided by social norms requires circuits that guide first-person behaviour to have access to information in circuits specialised for processing social information.

The paracingulate cortex is well known for its role in processing social information and particularly the mental states of others (Frith and Frith, 2003; Amodio and Frith, 2006; Frith and Frith, 2006). Recently, two studies have suggested that activity in this area may vary parametrically with the value of reward-related predictions during social interactions. Hampton et al. (2008) showed that activity in this area was scaled by the expected reward value on each trial. This value signal incorporated the influence that the subject's choice would have on their behavioural strategy. Behrens et al., (2008) also showed that activity in this area

was scaled by the probability that another's advice about rewarding values was truthful. Thus, activity in this area appears to be scaled by values that guide decision-making when these values are related to the mental states of others. The results in this chapter support the claims made in these two studies, showing that the paracingulate cortex is engaged when processing other's valuations of a delayed financial reward. However, unlike these two studies, this study shows that the paracingulate cortex is not only sensitive to the valuation of one other individual during social interactions. Here, the paracingulate cortex was shown to be engaged by social norm valuations, that reflect the value assigned to an action by the majority of others.

This study is not the first to investigate the neural circuitry which underpins the processing of social norms. Several studies have looked at how attitudes are processed when deciding whether to punish others for violations of social norms (Buckholtz et al., 2008). Other studies have also investigated the neural processes that underpin decisions of whether to conform to a norm (Berns et al., 2005; Spitzer et al., 2007; Klucharev et al., 2009) and also accidental violations of social norms (Berthoz et al., 2002; Berthoz et al., 2006; Spitzer et al., 2007; Prehn et al., 2008). None, however, have examined the processing of rewards temporally discounted in a manner that conforms to a social norm. Intriguingly, one previous study found activity in the ACC when processing normative valuations of facial attractiveness (Klucharev et al., 2009). Specifically they found that the RCZ coded for the discrepancy between a subjects' judgement of the attractiveness of another's face and its normative attractiveness. The magnitude of the signal in the RCZ correlated with the extent to which subjects conformed to the social norm on future trials. This suggests that first-person valuations are updated due to the discrepancy with a social norm valuation, much like a prediction error signal. In this study, activity in the ACCs signalled both subjective and normative values, suggesting that this region codes values that guide actions, regardless of the source of the valuations. Thus, both studies suggest that the ACCs may be important for processing normative values that guide first-person actions. However, this study had the advantage of testing whether computational models that have been shown to underpin first-person choice behaviour, can also explain behaviours when they are dictated by a social norm valuation.

6.5.3 Caveats and Limitations

An important feature of the design of this study was that subjects performed their own preferences on some trials and indicated a normative preference on other trials. The subjects were therefore making real economic decisions about the payment they would receive for their participation in the experiment. However, subjects were told prior to the experiment that they would not be paid based on their performance of the normative trials. This was necessary in order to maintain experimental control. If subjects had been paid based on the normative trials, they may have biased their decisions on those trials towards their own subjective preferences. As a result, subjects' performance on the social norm trials would have been different from the actual norm and also would have reflected a subjective valuation and not a normative valuation. In addition, it would have been unethical to pay participants based on their choices on the normative trials, when the normative valuations learnt during training were in fact fictional (there is no normative behaviour on this task). It is therefore plausible that any differences in activity between the subjective and normative conditions could be explained by a difference in incentive for performing each trial type. However, recently Bickel et al. (Bickel et al., 2009) used fMRI to examine activity in the brains of subjects performing an intertemporal choice task, for either real or fictive financial rewards. The results showed activity in the ACCs, ventral striatum and PCC in both the fictive and real monetary conditions. Crucially, no areas of the brain showed a significant difference in activity between the real or fictive subjective valuations. This result shows that areas of the brain that are involved in delay-discounting respond regardless of whether choices are made between real monetary values or not. In addition, behavioural studies have shown that when using real or fictitious rewards, behaviour can be equally well explained by a hyperbolic function, as was also the case in this study (Johnson and Bickel, 2002; Madden et al., 2003). It therefore follows that differences in brain activity between the normative trials (where the rewards were hypothetical) and subjective trials (where a real financial reward was available) in this chapter, cannot be accounted for by differences in the incentive available to participants on each type of trial.

The design of this study was such that subjects learnt social norm preferences that were actually fictional. These normative preferences were the output of a hyperbolic preference function. This model was chosen due to the substantial body of evidence that highlights this model as fitting closely with the behaviour of subjects on delay-discounting tasks (Rachlin et al., 1991; Green et al., 1994a; Kirby and Marakovic, 1996; Green et al., 1997; Kirby, 1997;

Mazur, 1997; Richards et al., 1997; Mazur, 2001; Madden et al., 2003). However, there is considerable debate as to the best model for explaining temporal discounting behaviours (Green et al., 1994a). Whilst there is evidence that hyperbolic functions are better at predicting temporal discounting behaviours than exponential functions (Green et al., 1997; Kirby, 1997; Johnson and Bickel, 2002; Madden et al., 2003; Kable and Glimcher, 2007), there are other models for which the hyperbolic function does not provide a better fit to the data. Area under the curve methods provide a univariate approach to examining discounting curves which perform as well as a hyperbolic function (Myerson et al., 2001). Utility based models that separate out a discounting into a hyperbolic discounting component and utility valuation component have been shown to predict behaviour better than a hyperbolic function on its own (Pine et al., 2009). Pine et al., (2009) examined the neural correlates of such a utility based model. In their study, subjects performed a delay-discounting task where they were required to choose between two delayed options, rather than between a delayed and an immediate option. Unlike in this study, or that of Kable and Glimcher (2007), they found that activity in the ACCs signalled the degree of discrepancy between the value of the two delayed options. One could argue, therefore, that the reason for activity being found in the ACCs to process subjective value is because a model that does not best explain the behavioural data was used. However, it should be noted that Pine et al. (2009) found that activity in the ACCs was modulated by subjective values, but it was the value of both the delayed options that modulated activity in this area. Thus, the results of Pine et al. (2009) still support the claims in this chapter, that activity in the ACCs codes for subjective and normative valuations of delayed rewards.

Whilst this chapter has interpreted the results as reflecting the neural processes driving normative behaviours, it is important to note that an alternative interpretation. Specifically subjects could have performed this task by simply retrieving valuations of delayed rewards from memory on the normative trials. This account is given credence by the design of the task, as the delayed options were the same in the training session when they learnt the normative behaviour, and the scanning session when they reproduced the normative behaviour. Subjects could therefore simply retrieve the behavioural responses from the training sessions to guide their behaviour on the normative trials in the scanning session. However, whilst this offers a plausible alternative interpretation there are two important reasons why the results of this study being reflective of the processing of social norms, rather than memory retrieval, is a more parsimonious account. Firstly, there is little evidence of the portions of the paracingulate cortex or the ACCg that were activated in this study being engaged when storing or retrieving

information from memory (Bush et al., 2000). Lesions to the ACCg in the nonhuman primate do not disturb tasks that require information to be retrieved from memory (Rudebeck et al., 2006) and neuroimaging studies do not typically report differential activity in these areas during the processing of semantic or procedural memories (Frith and Frith, 2003, Amodio and Frith, 2006). As stated in previous chapters, these areas are both implicated in the processing of social information. It therefore seems unlikely that activity in these particular regions is being driven by the retrieval of behaviours stored in memory. Secondly, activity that is reported in this study was found to vary parametrically with the value of a delayed reward and therefore was not simply differential activity between the delayed or immediate options, or remembering or not remembering the correct behaviour. In addition, the delayed reward options do not differ substantially in terms of the demands that they place on memory. As such, an account of the results as being related to the retrieval of memories, could not account for the parametric nature of activity in the areas that are reported here.

6.5.4 Summary

This study tested two hypotheses about the processing of delayed rewards in the ACC. The first hypothesis was that activity in the ACCs would signal the value of subjectively discounted rewards. In line with this hypothesis, activity in the ACCs varied with subjects own hyperbolically discounted rewards values at the time they made their responses. In addition, activity was found in the ventral striatum that varied with the subjectively discounted rewards values at the time the delayed option was presented. The second hypothesis was that activity in the in ACCg would signal the value of hyperbolically discounted delayed rewards that conformed to a social norm. In line with this hypothesis, activity in the ACCg varied with the normatively discounted rewards at the time that subjects made their response. In addition, activity in the paracingulate cortex varied with the value of the normatively discounted rewards at the time the delayed option was presented. These results suggest that making decisions that are guided by a social norm recruits parallel systems. One system is engaged in valuing different decision based on the social norm. Another system is engaged in processing values regardless of the source of the valuation. The exchange of information between these two systems may be crucial for making decisions which are guided by social norms.

Chapter 7: General Discussion

Studies investigating the functional properties of the ACC have implicated the region in a range of different processes including, pain (Rainville et al., 1997; Hutchison et al., 1999; Wager et al., 2004), response selection (Rushworth et al., 2004), the regulation of emotion (Bush et al., 2000; Phillips et al., 2003; Davis et al., 2005), error detection (Carter et al., 1998; Bush et al., 2000; Holroyd and Coles, 2002) and conflict monitoring (Botvinick et al., 1999; Barch et al., 2000; Kerns et al., 2004). The diversity of such processes has resulted in an absence of a clear picture of how this area contributes to behaviour. However, when one looks at the anatomy of the ACC, the reason that this region is implicated ubiquitously across a wide range of motor, cognitive and emotional processes becomes clear. Within the region commonly referred to as the ACC by functional imaging studies, there is considerable anatomical heterogeneity. Within this portion of the ACC there are at least 9 different cytoarchitectonic zones (Vogt et al., 1995; Bush et al., 2000; Paus, 2001; Palomero-Gallagher et al., 2008). In addition, the different sub-regions of the ACC are each connected to different areas of the brain (Pandya et al., 1981; Vogt and Pandya, 1987; Beckmann et al., 2009). Thus, an understanding of the ACC may come from localizing functions to specific sub-regions and testing hypotheses that are specific to each zone.

An aim of this thesis is to examine the functional properties of two different regions within the ACC. One region lies within the sulcal portion of the midcingulate cortex (MCC) which predominantly contains area 24c' (Vogt et al., 1995) and a second, is in the gyral MCC which predominantly contains areas 24a' and 24b' (Vogt et al., 1995). In order to discuss results consistently with the imaging literature and also due to the absence of any visible boundary between the MCC and its surrounding areas on an MRI image, these two areas have been referred to as the ACCs and the ACCg respectively throughout this thesis. However, it is important to note that the hypotheses tested and the overarching aim of the thesis has been to examine information processing in the MCC.

In chapter one, I highlighted how the absence of a theoretical account of the contribution of the ACCg to social cognition, was a result of the paucity of human functional imaging research which has tested the hypothesis that the ACCg processes social information. However, it was noted that a highly influential study made the claim that the ACCg processes similar information to the ACCs (Behrens et al., 2008). That study reported activity in the ACCg when subjects processed the outcomes of others' decisions and activity in the ACCs when subjects'

processed the outcomes of their own decisions. Therefore, as proposed in chapter one an understanding of the contribution of the ACCg to social cognition may come from drawing parallels with how information is processed in the ACCs.

There is considerable evidence that the ACCs is involved in processing information both at the time that choices are instructed and made, as well as at the at the time an outcome of a response is monitored (Rushworth et al., 2004; Walton et al., 2004; Kennerley et al., 2009). Such information processing conforms to the principles of a well founded theoretical framework, namely Reinforcement Learning Theory (RLT) (Rushworth et al., 2007; Holroyd and Coles, 2008; Rushworth and Behrens, 2008). There are two key principles to this theory (i) predictions are made about the value of an action and (ii) these values are updated when new information reveals that the prediction was erroneous. As outlined in chapter one, there is a considerable body of neurophysiology and neuroimaging evidence that supports the notion that the ACCs processes information in a manner that conforms to these two principles. In addition, a large number of studies have suggested that the predicted values encoded by neurons in the ACCs are modulated by variables that discount the value of rewards (Kennerley et al., 2009; Kennerley and Wallis, 2009; Hayden and Platt, 2010; Hillman and Bilkey, 2010). The four studies in this thesis tested whether the ACCg processes the discounted value of others' actions at the time that predictions are made and also whether the ACCg signals when others' predictions about the outcomes of their actions are erroneous. The results of all four studies can be summarised as follows:

- Chapter three: The ACCg responded exclusively when the outcome of a third-person's decision was unexpectedly positive. The ACCs responded to the unexpectedly positive outcomes of either a third-person's or a computer's responses.
- Chapter four: Activity in the ACCg varied with the discrepancy between a third-person's prediction and the actual outcomes of their action, in a manner that conformed to the computational principles of RLT.
- Chapter five: Activity in the ACCg varied with the net value of rewards, discounted by the effort expended by a third-person. Activity in the ACCs varied with the predicted level of effort regardless of whether the actions were to be performed by the subject or by a third-person.
- Chapter six: Activity in the ACCg varied with the value of delayed rewards, discounted in a manner that conformed to a social norm. Activity in the ACCs was found to vary with both the subjective and normative value of delayed rewards.

Broadly speaking, these results support the view that that the ACCs and the ACCg process variables that guide decision-making in a manner that conforms to RLT. In the ACCg this information is processed for third-person actions, whereas in the ACCs this information is processed regardless of whether it pertains to the actions of oneself or another.

7.1 Contributions of the ACCg to Social Cognition

One of the important features of this thesis was that specific anatomical predictions were made about the location of first-person and third-person information processing in the ACC. In this section, I discuss the extent to which the results support the claim that the gyral surface of the MCC is engaged when processing information about others' actions. Specifically, this section examines whether the ACCg processes the discounted value of others' actions and also whether information processing in this area conforms to the principles of RLT.

7.1.1 Does a portion of the ACC respond exclusively to social information?

All four experimental chapters in this thesis reported activity in the ACCg when processing social information. Chapters' three and four reported activity in this region when others' predictions were erroneous and chapters five and six reported activity in this area which varied with others' valuations of actions. Importantly, the ACCg was not activated in any of the studies by first-person predictions or when first-person predictions were erroneous. The notion that it is the ACCg and not the ACCs that processes social information was first proposed by Rudebeck et al. (2006a). As stated in chapter one, they examined the time between the presentation of a food item and a monkey reaching for the food item, when the food was presented simultaneously with an additional stimulus. When this stimulus was that of another monkey, or a human, there was latency before the monkey reached for the food item. This latency was reduced when the food item when was presented simultaneously with a neutral stimulus (such as moving dots). Lesions to the ACCg, but not the ACCs or the OFC, resulted in a significant decrease in the latency before the monkey reached for the food when it was paired with a social stimulus (see fig.1.4 in chapter one). The ACCg lesions also resulted in changes in the monkey's behaviour when interacting with other conspecifics, which did not occur following OFC or ACCs lesions. Thus, lesions to the ACCg and not the ACCs impact upon the processing of social information. The results of Rudebeck et al. (2006a) compliment the results in this thesis in highlighting the ACCg as important for processing social information.

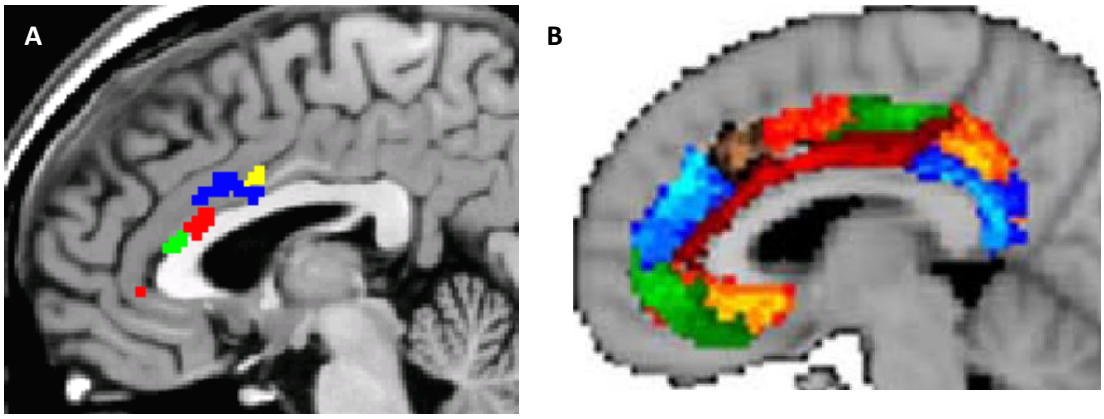

Fig. 7.1 (A) Activity shown in the ACCg for the four studies reported in the thesis, shown on the MNI template. Results from chapter three (yellow), in which the ACCg signalled when a third-person would receive an unexpectedly positive outcome; chapter four (green), in which activity in the ACCg signalled the discrepancy between a third-person's prediction and the actual outcome known by the subject; chapter five (red) in which activity in the ACCg varied with net value of a third-person's effortful actions and chapter six (blue) in which activity in the ACCg varied with the normative value of a delayed reward. (B) Results of DWI based parcellation of the cingulate cortex by Beckmann et al. (2009). The results in the studies of this thesis all putatively fall within the ACCg cluster shown in red. (C) An illustrative schematic of the cytoarchitectural zones of the ACCg taken from Bush et al (2000). Putatively, each of the results in (A) falls within area 24b'.

The results in this thesis all fell on the gyral surface of the MCC. Rudebeck et al. (2006a), lesioned the entirety of areas 24a and 24b and the majority of areas 24a' and 24b' on the gyral surface of the cingulate cortex in the macaques' brains. However, the most posterior portions of areas 24a' and 24b' were not lesioned. It should be noted here, that the same criteria were used to define these regions in the macaque and the human brain by Vogt and his colleagues (Vogt et al., 1987; Vogt et al., 1995). However, although the regions are homologous, it is difficult to determine whether the activations reported in this thesis fall within the same portion of the MCC that was lesioned. Nevertheless, figure 7.1 highlights how the clusters in two of the studies in this thesis, putatively lie in the most rostral portion of area 24a' and 24b', and therefore are likely to fall within the same portion of the MCC which was lesioned by Rudebeck et al., (2006a) in the macaque monkeys. The clusters activated in chapter three and five lie in the posterior portions of the MCC, and it is therefore not clear whether these fall in

the portion of the gyrus that was lesioned. The study by Behrens et al., (2008) that has been discussed previously in this thesis, reported activity in the ACCg that varied with the volatility of the advice provided by a third-person at the time of the outcomes of the subjects' decisions. Interestingly, the cluster that responded to the third-person information (the volatility of the confederate advice) in their study fell within a slightly more posterior portions of the gyral MCC, possibly in the area that was not lesioned by Rudebeck et al., (2006a). This cluster overlaps with the results of the two chapters which showed activity in posterior MCC. Another fMRI study by Apps et al. (accepted), also reported activity in the ACCg when processing the outcomes of others' decisions. In their task, subjects learnt a series of arbitrary associations between instruction cues and actions, and also observed different arbitrary associations being learnt by a third-person and a computer. Activity was examined at the time of the outcomes of each agent's decisions. There could be one of three possible outcomes on each trial: correct, incorrect, or ambiguous (a neutral stimulus that was uninformative about the accuracy of the response). They reported that activity in the ACCg, also in a more posterior portion of the MCC, was increased when the outcome of the third-person's trial was unexpectedly ambiguous. This region showed no response to any of the outcomes of the first-person's or computer's trials. The results of previous research and those in this thesis therefore support the notion that the overarching property of the gyrus in the MCC is to process social information.

Examination of the anatomical location of the results from each of the four studies in this thesis shows that the activated clusters all fall into an anatomical area that lies on the gyral surface of the cingulate cortex. Specifically the clusters lie posterior to the genus of the corpus callosum, inferior to the sulcus and anterior to the PCC. This anatomical area corresponds to the gyral surface of the MCC and therefore suggests that each of the activations reported in the ACCg in this thesis fall within either areas 24a' or 24b'. Figure 7.1 illustrates the correspondence between the locations of these clusters and areas 24a' and 24b' in the MCC. Interestingly, Beckmann et al. (2009) showed that a portion of the gyral surface of the cingulate cortex, extending over the whole of the MCC area, may have different connectional properties to the rest of the cingulate cortex. They parcellated the cingulate cortex using Diffusion-Weighted Imaging (DWI), showing that the gyral MCC shows a different pattern of connectivity from other portions of the cingulate cortex. An examination of figure 7.1 shows that the activations in the ACCg that are reported in this thesis all appear to overlap with the cluster resulting from the DWI paracellation, suggesting that the results in this thesis all fall within an area that has distinct connections from the other portions of the cingulate cortex.

The results of tracer studies in nonhuman primates support the notion that the connectivity profile of the gyral surface of the MCC is distinct from any other portion of the cingulate cortex. The most notable distinction is that this area is connected to posterior portions of the superior temporal sulcus, the paracingulate cortex and the temporal poles that are not found in the adjacent sulcus (Pandya et al., 1981; Vogt and Pandya, 1987; Seltzer and Pandya, 1989). These areas comprise the core-circuit which is engaged when mentalizing (Frith and Frith, 2003). Thus, the clusters activated in each of the studies in this thesis all fall within a portion of the MCC that has a connectivity profile that is distinct from other regions of the cingulate cortex. Specifically, the clusters fall within the only portion of the cingulate cortex that has access to a network that is specialised for processing social information.

An important finding in three of the chapters in this thesis is that the ACCs responds to the predicted values or erroneous predictions of others. In chapter five, activity in the ACCs varied with the level of effort on the third-person's trials and in chapter six, activity in the ACCs covaried with normative valuations of delayed rewards. In addition, in chapter three activity in the ACCs was found when the outcome of a third-person's response was unexpectedly positive. It could therefore be argued that the ACCs, like the ACCg, processes others' predictions and outcomes, invalidating the claim that it is the gyral surface of the MCC and not the sulcus that processes social information. However, crucially in none of the studies in this thesis was activity found in the ACCs exclusively when subjects' processed the predicted values, or the erroneous predictions of others. In each study, activity in the ACCs was also found to process first-person predicted values or first-person prediction errors (apart from chapter 4 where there was no first-person predicted value or prediction error related event) Thus, the ACCs is not engaged exclusively when processing social information, suggesting that it may process the predicted value of actions and signal erroneous predictions about actions, regardless of the source of the decision that led to the prediction. This therefore supports the claim that the ACCg is the only region of the cingulate cortex that is involved in processing the predictions and erroneous predictions of others

One caveat to this interpretation of ACCg function is that very few studies have tested specifically whether it is the gyrus, and not the sulcus, that is engaged in processing social information. As outlined in chapter one, in general the large corpus of studies investigating the neural antecedents of mentalizing abilities have not reported activity in the ACCg when subjects are processing the mental states of others (Van Overwalle and Baetens, 2009). In that chapter the methodological and conceptual reasons for why these studies may not have detected activity in the ACCg when subjects were processing information about others were

outlined. Typically, with an absence of neuroimaging data to support one's claims, one would examine the neurophysiological properties of neurons in the area under question. However, as also outlined in chapter one, to the best of my knowledge there has been no study that has recorded from neurons in the ACCg when subjects are processing a social stimulus. Thus, whilst the studies in this thesis were consistent in finding activity in the ACCg when processing other's predictions and erroneous predictions, there are only a limited number of other studies that support its claims.

7.1.2 Does the ACCg process the variables that guide other's decision-making?

In the ACCs, single-unit recording studies have shown that the there are neurons in which the activity varies with reward magnitude, reward probability and effort-related costs (Kennerley et al., 2009; Kennerley and Wallis, 2009; Hayden and Platt, 2010). Neuroimaging studies have also shown that activity in the sulcus of the MCC is modulated by reward probability and magnitude (Knutson et al., 2000; Knutson and Cooper, 2005; Knutson and Bossaerts, 2007). Activity in this area is also modulated by variables that discount rewards such as temporal delays before receipt (Kable and Glimcher, 2007; Luhmann et al., 2008; Bickel et al., 2009; Peters and Buchel, 2009) and the amount of effort to be expended (Croxson et al., 2009). This thesis examined whether the ACCg processes these variables about others' decisions, specifically examining whether the ACCg processes the net value of others' actions. Chapter five reported that activity in two different portions of the ACCg was engaged when subjects were processing the effort-discounted value or the undiscounted value of others' rewards. In chapter six, activity was also reported in the ACCg that signalled the value of delayed rewards that were discounted in a manner that conformed to a social norm. These two chapters therefore support the claims that the ACCg processes the effort-discounted and delay-discounted value of others' rewards, in the same manner that the ACCs processes the value of one's own rewards. However, the ACCg processes this information about others' valuations of rewards. None of the studies in this thesis explicitly examined the processing of reward probability. However, both chapter three and four showed that the ACCg signals the discrepancy between another's prediction and the actual outcome, when the predictions were probabilistic in nature. As such, the third-person's prediction, and therefore the error, was also probabilistic. Overall, the results of this thesis highlight the ACCg as a region which codes the

same variables that are coded in the ACCs. However, the evidence in this thesis suggests that the ACCg processes this information about the decisions and actions of others.

Anatomical evidence indicates that the ACCg may have access to information in areas that are engaged when subjects' discount the value of rewards. The gyral MCC has connections to the core of the nucleus accumbens and the CMAr (Dum and Strick, 1991; Kunishio and Haber, 1994; Morecraft and Van Hoesen, 1998), as well as portions of the ACCs that lie rostral to it (Pandya et al., 1981; Vogt and Pandya, 1987; Morecraft et al., 1992). Both the nucleus accumbens and the CMAr are implicated in processing effort-discounted reward values. Previous neuroimaging studies and chapter five in this thesis have reported activity in the CMAr and the nucleus accumbens that varies with the value of effort-discounted rewards (Botvinick et al., 2009; Croxson et al., 2009). In addition, neurons in the CMAr have been found to increase their firing rate with the predicted level of effort to be expended (Kennerley et al., 2009) and lesions to the nucleus accumbens disrupt choices between rewarding options with different levels of effort required for their receipt (Ishiwari et al., 2004; Salamone et al., 2007). As such, the gyral surface of the MCC would appear to have access to information about the effort-discounted value of rewards. Similarly, the MCC is connected to areas of the brain that have been found to be engaged when subjects discount the value of delayed rewards, including the intraparietal sulcus the PCC, as well as the nucleus accumbens and more anterior portions of the MCC (Pandya et al., 1981; Hurleygius and Neafsey, 1986; Vogt and Pandya, 1987; Buckwalter et al., 2002; Vianna et al., 2002; Kable and Glimcher, 2010; Louie and Glimcher, 2010; Prevost et al., 2010). Thus, the ACCg may have access to information in areas that process both the effort-discounted and the temporally discounted value of rewards. These connections and those to the core-circuit engaged when mentalizing discussed earlier, highlight the ACCg as a candidate for processing the discounted value of others' rewards.

Thus far, it has been claimed that the ACCg is engaged when the reward value associated with others' actions are processed. However, as outlined in chapter one, the conflict monitoring account of the functional properties of the ACC is the most prominent in the literature (Carter et al., 1998; Botvinick, 2007; Cole et al., 2009, 2010). Previously in this thesis, it was argued that the conflict monitoring perspective could not account for the results of a number of studies investigating the functional properties of the ACC. However, given the prominence of this theory in the literature, it is pertinent to discuss whether the results reported in this thesis could be accounted for by its predictions. In chapter three, the subjects' task was to determine whether the outcome of a third-person's or computer's response was the same as that which was predicted on the basis of a fixed ratio of probabilities. Specifically, on 1/3 of trials the

outcome was different from that which was predicted. The subjects indicated on each trial whether the outcome was discrepant from the predicted outcome by making one of two responses on a keypad. It could be argued that on the trials where there was a discrepancy between a prediction and an outcome, there was conflict between the subject's prepotent response (i.e. a button press indicating that there was no discrepancy) and the response that was now required (i.e. a button press indicating that there is a discrepancy). However, in this study, activity in both the ACCs and the ACCg was found when an outcome was unexpectedly positive and not when the outcome was unexpectedly negative. Both of these trial types would require the cancellation of the same prepotent response. Thus, the absence of a response in the ACC to unexpectedly negative outcomes suggests that changes in activity in this area could not be explained by changes in the level of conflict.

In chapter five, activity was examined at the time of instruction cues that signalled the level of reward and the amount of effort to be expended. It has been argued that when choices are made between different options, the ACCs processes how much conflict there is between them (Botvinick, 2007). However, in chapter five, subjects were not presented with a choice, but were instructed how much effort needed to be expended to receive a reward. Thus, the identified responses in the ACCs and ACCg could not be explained by conflicts between different options.

Finally, in chapter six, activity in the ACC varied with normative and subjective valuations of delayed rewards. It could be argued that the activity in the ACC reflected the degree of conflict between the delayed and immediate options. However, activity in both the ACCg and the ACCs increased as the value of the discounted rewards increased. Thus, as there was an increase in conflict between the two options, there was a decrease in the ACC response. Such a profile of activity contradicts with the conflict monitoring account, which would predict that the closer the values of the delayed and immediate options are, the greater the conflict between the options and therefore the greater the response would be in the ACC. Thus, the responses in the ACC were diametrically opposite to those which would be predicted by the conflict monitoring account. Therefore, in this thesis, there is little evidence in this thesis to support the claims of the conflict monitoring account of ACC function.

7.1.3 Does information processing in the ACCg conform to the principles of RLT?

In chapter one, I suggested that the ACCs may process information in a manner that conforms to the principles of RLT. RLT posits that there are two main processes that underpin learning and decision-making (Sutton and Barto, 1998; Dayan and Balleine, 2002; Schultz, 2006; Dayan and Daw, 2008). Firstly, predictions are made about the value of alternative choices. These predicted values are used to guide decision-making, and once a choice has been made, the values are coded as the predicted outcome of the chosen response. The second component of RLT is the prediction error signal that updates the value of an action when an outcome was unexpected. To examine whether information processing in the ACCg conformed to the principles of RLT, it was therefore important to examine activity at the time of cues that (i) signalled the predicted value of others' actions and (ii) signalled a discrepancy between the predicted value of another and their actual outcome (the prediction error component). The results of this thesis support the notion that the ACCg processes information in a manner that conforms to the principles of RLT. In chapters' three and four, activity in the ACCg signalled the discrepancy between another's prediction about the value of the actions and the actual value of that action known by the subject. These two studies therefore support the notion that the ACCg processes the prediction error component of RLT. In addition, in chapter five activity in the ACCg varied with the net value of another's effort-discounted rewards at the time of instruction cues, suggesting that the ACCg processes the predicted value of others' actions. In addition, in chapter six activity in the ACCg varied with a normative valuation of discounted rewards at the time of cues that instructed the choice to be made. Although this study did not examine whether the ACCg processed another's predictions *per se*., as it was the value according to a social norm that was being processed, it still highlights that the ACCg processes others' valuations of rewards at the time when predictions can be made. Thus, the results in this thesis broadly support the notion the ACCg processes the predicted values and erroneous predictions of others in a manner that conforms to RLT.

To the best of my knowledge, the research in this thesis is the first to test the hypothesis that activity in the ACCg varies with the predicted value of others' rewards at the time that another is instructed about their actions. However, one study has shown that different portions of the ACC are engaged when processing other's actions from those that are engaged when making a choice oneself. An influential study by Tomlin et al., (2006) used a pattern analysis approach to examine activity in the ACC during a game-theoretic trust task. In this game, the subject and a

partner performed an iterative task where one player is required to invest money, which is tripled for the other player who makes a decision as to how much money to repay the investor (Tomlin et al., 2006). Tomlin et al., (2006) found that the pattern of activity in the MCC was different when subjects invested money, from when they monitored the decisions of the other player, i.e. at the point in time when they could infer the value that the other was placing on the rewards. Unfortunately, due to the pattern analysis technique employed, it is not possible to comment on whether this was a result of specific changes in the gyral MCC and not the sulcal MCC. However, this study offers tentative support for the notion that distinct portions of the ACC process one's own and other's predictions about the value of their actions.

Whilst few studies examined the role of the ACCg in processing social information, two studies have reported activity in the ACCg when a subject monitors the outcomes of others' decisions. In one study, subjects were required to perform, or observe a third-person performing, a go/no-go task. They reported that the gyral surface of the MCC, was activated exclusively to the erroneous inhibitions and erroneous responses of the third-person (Shane et al., 2008). This supports the notion that the ACCg signals when others' responses are erroneous at the time that the outcome of their response is revealed. Another study by Apps et al., (accepted), found that a portion of the ACCg responded exclusively when the prediction that the outcome of a third-person's decision would be informative was erroneous. Thus, these results tentatively support the conclusions that can be drawn from the studies in this thesis, which the ACCg signals when the predictions of others are erroneous.

The gyrus of the MCC is connected to areas that that signal reward prediction errors, supporting the notion that the ACCg may also code prediction error signals. The anterior portions of the MCC, on the gyral surface, receives monosynaptic projections from neurons in the Ventral Tegmental Area (VTA) in the midbrain (Williams and Goldman-Rakic, 1998). Neurons in the VTA increase their firing rate quantitatively with the magnitude of the discrepancy between a prediction and error (Hollerman and Schultz, 1998; Schultz, 1998; D'Ardenne et al., 2008; Bayer and Glimcher, 2005). As yet, unfortunately, no study has examined whether such connections are also found from the posterior portions of the MCC. In addition, both the posterior and anterior portions of the MCC have connections to the CMAr which, as discussed throughout the thesis, contains neurons which signal when predictions are erroneous (Matsumoto et al., 2007; Hayden et al., 2011b). In addition, the ACCg also has connections to areas that functional imaging research in humans has found to be activated when processing the discrepancy between predictions about rewarding outcomes and the actual outcome. This includes both the OFC (Carmichael and Price, 1995; Ramnani et al., 2004)

and the nucleus accumbens (Kunishio and Haber, 1994; Schultz and Dickinson, 2000; Brovelli et al., 2008; D'Ardenne et al., 2008). The ACCg therefore has access to information in many regions which signal when predictions are erroneous, highlighting it as a candidate region for processing prediction errors as well.

The results in this thesis support the claim that information processing in the ACCg conforms to the principles of RLT. However, the nature of the prediction error signals in the ACCg and the ACCs are clearly distinct. Specifically, the ACCs signals when first-person predictions about the outcome of a decision is discrepant from its actual outcome (Frank et al., 2005; Matsumoto et al., 2007; Holroyd and Coles, 2008; Jocham et al., 2009). In contrast, the ACCg signals the discrepancy between the prediction of a third-person and the actual outcome which is known by the first-person and not by the third-person. Therefore, the studies in this thesis examined activity in the ACCg when the subject processed the erroneous predictions of a third-person and not when they were monitoring and processing another's prediction error, i.e. activity was not examined at the point in time when an outcome revealed to the third-person that their prediction was erroneous. The error signals identified in this thesis therefore do not conform to a strict interpretation of RLT, which states that a prediction error signal updates predicted values (Sutton and Barto, 1981; Sutton and Barto, 1998; Dayan and Daw, 2008). In the ACCg the prediction signals occurred when only the first-person was able to process the discrepancy between the third-persons' prediction and the actual outcome. The third-person could not have been updating their valuation of an action at that time, as they still had not been informed of the outcome of their action. Thus, the identified prediction error signal does not reflect the update signal of another *per se*. However, it should be noted that in both of the studies where subjects' processed the discrepancy between a third-person's predictions and their actual outcomes, the outcomes were also revealed to the third-person in a subsequent trial event. Therefore, one possible explanation of the ACCg prediction error signal is that the subjects were processing the update of the third-persons predictions when an event revealed the outcome to them, knowing that the third-person would update their valuation at a subsequent trial event. Unfortunately, the designs of these studies did not permit a specific and stringent test of this hypothesis.

7.2 ACCs: Signalling the value of actions during decision-making

In chapter one, I outlined a framework of how the ACCs might contribute to the processing of actions and also to the monitoring of outcomes. An aim of this thesis was to apply the same framework to the ACCg, in order to examine its contribution to social cognition. However, it is also pertinent to discuss whether the studies in this thesis support the notion that the ACCs processes the value of actions and that information processing in this area conforms to the principles of RLT.

Broadly speaking, each of the studies reported in this thesis supports the claims made about information processing in the ACCs in chapter one. As discussed earlier, in chapter three the ACCs was activated whenever there was an unexpectedly positive outcome of a response by either a third-person or a computer, suggesting that the ACCs was coding for the subject's own erroneous predictions. This therefore supports the substantial body of neurophysiology research that has found neurons in the CMAr and also rostral portions of the sulcal MCC, in which the spike frequency varies with the degree of the discrepancy between a prediction and an outcome in the manner predicted by RLT (Amiez et al., 2005; Matsumoto et al., 2007; Sallet et al., 2007; Quilodran et al., 2008). In addition, in chapters' five and six activity in the ACCs varied with predictions about the value of rewards at the time that actions were instructed or made. This supports the findings of several neurophysiology and neuroimaging studies that report activity in the ACCs that increases parametrically with the predicted value of rewards (Knutson et al., 2000; Shidara and Richmond, 2002; Rogers et al., 2004; Williams et al., 2004; Amiez et al., 2005; Quilodran et al., 2008; Rolls et al., 2008; Kennerley et al., 2009; Hayden and Platt, 2010). In addition, these two chapters also support the claim that the ACCs is engaged when subjects' discount the value of rewards. Chapter five showed that activity in the ACCs varied with the level of effort that was to be expended by a subject, in line with previous neurophysiology studies that have found neurons in this area in which the spike frequency increases with the number of actions that are to be performed (Kennerley et al., 2009; Kennerley and Wallis, 2009). Finally, chapter six found that activity in two portions of the ACCs varied with the temporally discounted values of rewards, replicating the findings of several previous studies (Kable and Glimcher, 2007; Peters and Buchel, 2009; Kable and Glimcher, 2010). Thus, the results of this thesis corroborate previous findings that the ACCs processes predictions about the discounted value of actions and also signals when these are erroneous, in the manner predicted by RLT.

Throughout this thesis I have suggested that variables which discount the value of rewards are processed in the ACCs. Specifically, I have suggested that these variables are processed in caudal portions of the sulcal MCC (the RCZ in humans). However, at this juncture it is important to note that several areas, including the OFC, the Ventromedial Prefrontal Cortex (VmPFC), the ventral striatum, the PCC, and portions of the intraparietal cortex have all been implicated in processing similar reward-related values, both in this thesis and in other studies (Kennerley et al., 2009; Rushworth et al., 2009; Louie and Glimcher, 2010; Pearson et al., 2011). Each of these areas contains neurons which respond parametrically to value information during decision-making, although they may do so during different behaviours and during different events in a task. If the RCZ codes the same values as other areas, what specific contribution does it make to guiding behaviour? To answer this question, one has to look at the differences between the connectional properties of each of these areas. As outlined in chapter one, the rostral CMA (CMAr), which lies in the sulcal MCC (corresponding to a portion of the RCZ in the human) has strong connections to areas of the brain which process rewards (Vogt and Pandya, 1987; Carmichael and Price, 1995) and also to areas in the motor system (Dum and Strick, 1991; Luppino et al., 1991; Devinsky et al., 1995). In particular, this region has connections directly to the spinal cord, to primary motor cortex, to premotor cortex and also to the supplementary (SMA) and pre-supplementary motor areas (pre-SMA) (Wang et al., 2001). In the most rostral portions there are fewer neurons that project to the primary and premotor cortices, although connections to the SMA and pre-SMA are found up to the MCC-ACC border (Wang et al., 2001). Crucially, neurons that are sensitive to variables that guide decision-making are found across the length of the MCC (Kennerley et al., 2009; Kennerley and Wallis, 2009; Hayden and Platt, 2010; Hillman and Bilkey, 2010; Hayden et al., 2011b; Hayden et al., 2011a). This would suggest that reward-related processing in the ACCs becomes increasingly abstract rostrally along the rostral-caudal axis, consistent with the general organisation of the frontal lobe. In the CMAr, the information processed is related specifically to providing motivation for the performance of an action. In contrast, information processing in the more anterior portions may process reward value in a more abstract form.

There is some evidence in this thesis for the posterior MCC processing the value of actions. In chapters three and five, activity was found in the posterior MCC/RCZ, when an unexpected outcome of a response occurred and when processing the number of actions that were to be performed respectively. In chapter six, activity in this area varied with the temporally discounted value of rewards at the time that an action was selected. A portion of the anterior MCC also varied with the temporally discounted value of rewards, however, it processed this

information at the time that the delayed option was presented and not at the time the action was selected. This tentatively supports the notion that the overarching property of the sulcal MCC is sensitivity to reward value. In the anterior portions of this area, the value processed may be abstract in relation to the action, whereas in the posterior portions the reward information is associated specifically with an action.

The notion that the ACCs is important for associating rewards with actions is also supported by the effects of lesions on tasks that require associations to be made between actions and rewards. Kennerley et al., (2006; and later Rudebeck et al., 2008 who analysed the data using a different approach) reported two experiments which were performed on monkeys who had lesions extending along the sulcus from the most rostral portions of the sulcal ACC to the most caudal portions of the MCC. In the first experiment, monkeys performed a task where one of two movements of a joystick would result in a rewarding outcome. The rewards were deterministic, such that monkeys had to sustain a rewarded movement for 25 trials, after which the action-outcome contingencies were switched, such that the previously unrewarded action would now be rewarded and the other would not. Monkeys in a control group and the ACCs lesioned monkeys were able to switch to the alternative action following an error, i.e. the trial immediately following the switch in action-outcome contingencies. However, the lesioned monkeys, unlike the control group, were unable to sustain the performance of the rewarded action following a switch, suggesting that the monkey was unable to associate the reward with the chosen response. In experiment two, monkeys performed similar a task, choosing between two lever movements. However, the reward associated with each action was not deterministic as in experiment one. Each action was assigned a probability of resulting in the reward that was in a fixed ratio with the probability that the other action would result in a reward. The reward ratios between options and the probability of each action being rewarded were stable in blocks of trials, such that subjects could learn them over several trials during a block. In separate blocks the ratios and probabilities were varied. Interestingly, the ACCs lesioned monkeys required significantly more trials to learn the ratios and the probabilities of each action, again suggesting that they were unable to associate actions with rewards (Rudebeck et al., 2008).

Further support for the claim that CMAr in the posterior MCC processes rewards is provided by the study Hayden et al., (2010). Hayden and Platt (Hayden and Platt, 2010) trained monkeys on a task where they were required to make one of seven different cued saccades to receive either a high or low reward. They found that neurons in this region were sensitive to the reward level, but only for a specific saccade direction. This argues against the notion that the

ACCs processes abstract information, suggesting that this area is engaged in assigning a value to a particular action. One could make the claim that this area is therefore processing attention related responses. However, Kennerley et al., (Kennerley and Wallis, 2009b) found that there were no neurons in the ACCs that coded for spatial location, suggesting that the neurons identified in the ACCs by Hayden and Platt (2010) were firing in response to the specific action (the saccade), and not the location of the target. This therefore suggests that neurons in the ACCs associate a value with a specific action or series of movements. Thus, whilst many areas of the brain process values, in the ACCs these values may be specifically associated with actions.

At this juncture it is important to note that the MCC is implicated by a sizeable corpus of literature in the processing of noxious stimuli (Lenz et al., 1998; Hutchison et al., 1999; Becerra et al., 2001; Buchel et al., 2002). Single-cell recordings from neurons in both the posterior and anterior MCC in the human, have shown that there are neurons that increase their spike frequency when thermal (both hot and cold) or mechanical painful stimulation is applied to the skin (Hutchison et al., 1999). A neuroimaging study that have applied noxious stimuli to the body reported activity in both the posterior and anterior portions of the MCC (Peyron et al., 2000). The location of activations in such studies has been shown to overlap considerably with the portions of the MCC that has been implicated in the processing of rewards, actions and conflict (Beckmann et al., 2009). This overlap could potentially be considered a limitation with the theoretical framework I have proposed.

Can the viewpoint that the MCC processes pain and the view that it processes the value of actions be reconciled? A recent model of the contribution of the MCC to pain processing shows a striking correspondence to the theoretical framework that has been outlined in this thesis (Shackman et al., 2011). This model suggested that that the function of this region may be to motivate behaviours to avoid painful stimulation. This opens up the possibility that pain, much like effort and reward probability, is a factor that modulates the value assigned to an action. The higher the level of pain experienced, the greater the motivation for acting to prevent continuous stimulation.

If the contribution of the MCC to pain processing is to motivate avoidance behaviour, it is important that the firing of neurons in this area is not related to the detection of pain. In accordance with this notion, electrical stimulation of neurons which respond to painful stimulation in the human MCC, does not result in subjects' reporting the detection of a painful stimulus (Hutchison et al., 1999). fMRI studies show that activity in the MCC is sensitive to

placebo-analgesics that reduce the perceived level of pain (Wager et al., 2004). Thus, the level of painful stimulation applied to the body is not directly related to activity in the MCC, suggesting that the MCC is not part of the system that detects painful stimulation, but another process that precedes or follows detection. Interestingly, a study which took single-unit recordings from neurons in the posterior portions of human MCC, found that some of the neurons in this region respond when painful stimulation can be predicted (Hutchison et al., 1999). These neurons responded before the painful stimulation occurred, when the subject was told they were about to receive a pinprick to their skin. This would therefore suggest that neurons in the MCC are engaged by the process of predicting the receipt of pain, before any pain could be detected. This could also be interpreted as being indicative of neurons coding the motivation for avoiding painful stimulation.

Single-unit recordings from the monkey MCC have found that some neurons in the MCC show increased firing rates on trials where monkeys perform a successful action to avoid painful stimulation (Iwata et al., 2005). These neurons are found in the same area that contains neurons which project to the spinal cord (Dum and Strick, 1996) and that contains neurons to which stimulation results in limb movements (Luppino et al., 1991). This might suggest that the ACCs signals the motivation for performing an action that avoids a painful stimulus. fMRI studies in humans also support this assertion. One study showed that activity in the MCC increases parametrically with the proximity of a tarantula (Mobbs et al., 2010), which could be indicative of an increased motivation for performing actions to avoid any pain that could be inflicted by the spider. In another study, activity in the MCC varied parametrically with the proximity of a virtual predator in a maze based hunting game, when the subject was acting as the prey (Mobbs et al., 2007). Increased activity was therefore related to an increased motivation for performing actions to avoid the predator.

Lesion studies support the notion that MCC is involved in processing pain, but contrast with the notion that it processes the motivation for avoidance behaviour. Lesions to the posterior portions of the MCC disrupt normal behavioural responses that occur when an animal is in a situation that could result in a painful outcome. In the study by Rudebeck et al. (2006a) that has already been described in detail, monkeys were presented with food at the same time as another stimulus. In one condition they presented the food simultaneously with a video of snake. In the control group, there was a considerable duration between the presentation of the stimuli and the monkey reaching for the food item. However, lesions to the ACCs resulted in a significantly reduced latency before the monkeys reached for the food. One interpretation of this is that the monkeys without lesions to the ACCs, discount the value of the food item

due to the perceived threat of the snake causing them pain. A lesion to the ACCs results in the value of the food not being discounted by the predicted level of pain associated with the threatening stimulus.

There are considerable parallels between how the ACCs processes pain and how the ACCs processes reward and effort. As stated previously in this chapter, Kennerley et al. (2009) showed how there are neurons in the ACCs that increase their firing rate as the reward magnitude (i.e. the motivation) associated with performing an action increases. However, there are also neurons that process the reward value discounted by the amount of effort (or cost). Thus, it is possible that pain may be processed in the ACCs in a similar manner to how rewards and effort are processed, providing motivation for actions and also discounting the value of rewards.

7.3 Theory of Mind, Simulation and Decision-making

As outlined in the introduction, historically there were two prominent accounts of the neural basis of social cognition. Simulation theory suggests that the cognitive mechanisms that underpin the processing of one's own actions and intentions are simulated when processing the same information about another (Gallese and Goldman, 1998; Gallese, 2007). The Theory of Mind account suggests that we understand others by making theories about their mental states (Frith and Frith, 2003). The discovery of mirror-neurons and the mirror-neuron system supported the claims of simulation theory, highlighting how areas of the brain that are engaged when performing an action are also engaged when observing another performing the same action (Gallese and Goldman, 1998; Rizzolatti and Craighero, 2004; Gallese, 2007). In addition, the discovery of a core-circuit that is engaged when processing others' mental states, added support for the theory of mind account (Frith and Frith, 1999, 2006). However, as discussed in chapter one, neither account is sufficient to explain all aspects of social cognition and neither can account for the fact that damage to other areas of the brain disrupts social behaviour. In this section, I argue that the principles on which these theories are based are still useful for understanding the neural basis of social cognition. However the theories need to be updated to reflect recent evidence about how the brain responds to rewards during social interactions.

Recently is has been suggested that the traditional account of the mirror-neuron system is not sufficient to explain how others' actions are predicted, or how the goals of their actions are understood (Kilner, 2011). Rather, such functions may be performed by other areas within the motor system, that also have mirror-like properties. Evidence in support of this notion has been provided by an fMRI study by Ramnani and Miall (2004). As outlined in previous chapters, they found activity in portions of the middle frontal gyrus (area 9/46) at the time of instruction cues, when subjects' could predict the actions of a third-person. Area 9/46 in the middle frontal gyrus is an area that is typically considered as important for processing abstract information about actions and is not considered part of the mirror-neuron system. This suggests that predicting which action another will perform, recruits parts of the motor system outside of the mirror-neuron circuit. Other studies have also found that parts of the motor system, not considered as parts of the mirror-neuron system, are activated when subjects observe others' actions. In an fMRI study on macaque monkeys, areas 45a, 45b and 46 on the lower bank of the principal sulcus showed increased activity when the monkeys observed others performing actions (Nelissen, 2005). These results therefore suggest that areas in the motor system, not considered a part of the mirror-neuron system, also have mirror-like properties.

In this thesis, there is some evidence that another part of the motor system, in the posterior portions of the MCC (the RCZ), an area that has strong connections to premotor area 44 (Picard and Strick, 1996), also has mirror-like properties. In chapters five and three, activity in this area processed predicted values and erroneous predictions respectively, regardless of whether these values were related to one's own or another's predictions. Similarly, in chapter six activity in the posterior MCC varied with both subjective and normative valuations of delayed rewards. These results suggest that this area codes information in a manner that is not agent-specific. Although these studies did not directly test a simulation theory hypothesis, they certainly provide evidence to support the notion that the RCZ, an area that is considered part of the motor system, also has mirror-like properties. Understanding others' actions may therefore require simulation in many different portions of the motor system.

The results of this thesis suggest that several other areas also have mirror-like properties. The insula and a portion of ventromedial area in the frontal lobe, at the juncture between the superior portions of the OFC and the inferior portions of the superior frontal gyrus on the medial, are well-known for processing reward values (Hare et al., 2009; Smith et al., 2010; Tricomi et al., 2010). However, in chapter four, activity in both of these areas varied with a third-person's prediction about the value of an action, suggesting that depending on the task

context, these areas may process values regardless of the source of the valuation. Chapters five and six also found activity in the VmPFC and anterior portions of the MCC (anterior to the portions described above) that varied with both one's own and others' valuations of effort and temporally discounted rewards. Two other studies that have been discussed in detail in this thesis (Behrens et al., 2008; Hampton et al., 2008), have shown that activity in the VmPFC varies with the valuation of a reward that combines first-person and third-person estimates of value. These studies therefore support the claim that this area may process the value of rewards, regardless of the agent who is valuing it. Interestingly, the studies in this thesis and previous neuroimaging studies that have already been discussed also show that these areas process prediction error signals (Burke et al., 2010; Preuschoff et al., 2008). Tentatively, therefore, I suggest that there are areas of the brain, including the VmPFC, the insula and the sulcus of the MCC, in which activity varies parametrically with statistical properties that are related to the value of rewards, regardless of whom will receive the reward. In these areas that have mirror-like properties, information is processed in a manner that conforms to RLT.

Whilst there is evidence that there are areas which are engaged when simulating the processing of others' actions and decisions, there is also evidence to suggest that there are areas that are recruited exclusively when processing social information. In chapter one, evidence was reported in support of the notion that there are a core-circuit of brain areas that are involved specifically in processing the mental states of others (Frith and Frith, 2003; Frith and Frith, 2006; Frith and Frith, 2010). In several chapters I have also discussed two studies, by Behrens et al., (2008) and by Hampton et al., (2008), that reported activity in the paracingulate cortex and the pSTS that varies parametrically with reward-related values during social interaction tasks. At several points I have also highlighted how these two previous neuroimaging studies have reported that areas in this network process prediction errors. In this thesis, I have reported evidence that suggests that the ACCg is similarly specialised, processing other's predictions and erroneous predictions about the value of their actions. These studies therefore provide evidence to suggest that the pSTS, the paracingulate cortex and the ACCg form an interconnected network of areas that process information about the mental states of others. These areas process this information in a manner that conforms to the principles of RLT

In summary, the results of this thesis support the notion that there is a network that is engaged exclusively when processing the decisions and actions of others. The ACCg falls within this network and may specifically be engaged by processing the value that others' place on their actions. However, this network is not sufficient for all aspects of social behaviour.

Interacting successfully in social environments may require others' value-related information to be simulated in networks which also process first-person reward valuations. These two systems, however, can be unified by the fact that they process information in a manner that conforms to the principles of RLT.

7.4 Future Directions

Like much scientific investigation, the work in this thesis has generated more questions than it has answered. Whilst it is not possible to report all of the possible extensions of this work here, I will try and summarise the three key areas in which the most interesting and relevant work could be pursued.

Perhaps the most important finding in this thesis is that it was specifically the gyral surface of the MCC that was involved in processing social information in each of the studies and not the sulcus. In chapter one, I presented a series of studies in which recordings were taken from the ACC during decision-making tasks. These studies showed that there are neurons in the ACC which are sensitive to variables which guide decision-making. However, in these studies electrodes were placed exclusively in the sulcus (Shidara and Richmond, 2002; Amiez et al., 2005; Sallet et al., 2007; Quilodran et al., 2008; Kennerley et al., 2009; Kennerley and Wallis, 2009) and to the best of my knowledge no previous study has recorded from neurons in the gyrus. As such, the studies in this thesis and one other functional imaging study by Behrens et al. (2008), are the only studies that have directly tested whether ACCg, and not the ACCs, is engaged when processing the value of others' actions in a manner that conforms to RLT. Thus, the most obvious extension to the research in this thesis would be to examine the functional properties of neurons in the ACCg. This research could be conducted in monkeys, or alternatively on patients who are about to receive a cingulotomy. To date, recordings of neurons in the gyrus have been absent from the literature, largely due to the practical issues that surround recording from a region in the depths of the medial wall. However, if such practical issues can be overcome, several questions based around the results of this thesis could be asked about what specific functional properties neurons in the ACCg have. Firstly, do neurons in the ACCg respond when monkeys observe the actions of another? Secondly, if neurons respond to others' actions, do they process the effort, reward magnitude and reward probability associated with others' actions, in the same manner as neurons in the ACCs do for

first-person actions? Thirdly, if neurons process the variables that discount the value of others' actions, are there neurons that multiplex this information, in the same manner as in the ACCs? Answering these questions may be difficult due to the impracticality of training monkeys on tasks where they process these sources of information about other agents. However, the answers to such questions may enable stronger conclusions to be drawn about the results of the studies reported in this thesis and also further the understanding of the contribution of the ACCg to social cognition.

In this thesis, I presented a framework for information processing in the ACCs, suggesting that the overarching function of the region is to process the value of actions during decision-making. Previous research and the results in this thesis, suggest that the ACC processes the value of rewards that are discounted by other variables such as effort and delay. However, I have also suggested that values are updated by prediction error signals that occur in the ACC. An important issue which has not yet been investigated is whether the ACC processes the discrepancy between the predicted net value of delay and effort-discounted rewards and their actual value i.e. do prediction error signals in the ACCs code for incidences when the level of effort expended is different from the expected level of effort? Further research investigating such prediction error signal, may therefore provide additional support for the notion that ACCs conforms to the principles of RLT.

In section 7.2 I suggested that pain may also be a variable that discounts the value of actions and, additionally may also be a motivation for the avoidance of actions. However, despite the evidence presented that tentatively supports this claim, there is an absence of neurophysiological or neuroimaging evidence that has directly tested whether such processes are localized to the ACCs. Therefore future research is required to investigate whether pain can be incorporated into the framework that has been presented in this thesis. Specific questions could be asked to tackle this issue: do neurons in the ACCs increase their firing rate in order to motivate a specific action to avoid a painful stimulus? Are there neurons in the ACCs that multiplex reward magnitude and pain, in a manner that discounts the value of a reward based on the amount of pain that will be suffered for its receipt? Finally, are there neurons in the ACCg that increase their firing rate exclusively when another is observed receiving pain, or performing an action to avoid a painful stimulus? Answers to these questions would provide considerable insight into how the ACC contributes to motivating useful actions and avoiding rewarding actions that are associated with significant costs.

References

Adolphs R (2002) Neural systems for recognizing emotion. Current Opinion in Neurobiology 12:169-177.

Aichhorn M, Perner J, Weiss B, Kronbichler M, Staffen W, Ladurner G (2009) Temporo-parietal Junction Activity in Theory-of-Mind Tasks: Falseness, Beliefs, or Attention. Journal of Cognitive Neuroscience 21:1179-1192.

Ainslie G (1975) Specious reward - behavioral theory of impulsiveness and impulse control. Psychological Bulletin 82:463-496.

Ainslie GW (1974) Impulse control in pigeons. Journal of the Experimental Analysis of Behavior 21:485-489.

Alexander GE, Crutcher MD (1990) Functional architecture of basal ganglia circuits - neural substrates of parallel processing. Trends in Neurosciences 13:266-271.

Alexander GE, Delong MR, Strick PL (1986) Parallel organization of functionally segregated circuits linking basal ganglia and cortex. Annual Review of Neuroscience 9:357-381.

Allman J, Hakeem A, Watson K (2002) Two phylogenetic specializations in the human brain. Neuroscientist 8:335-346.

Allman JM, Tetreault NA, Hakeem AY, Manaye KF, Semendeferi K, Erwin JM, Park S, Goubert V, Hof PR (2010) The von Economo neurons in frontoinsular and anterior cingulate cortex in great apes and humans. Brain Structure & Function 214:495-517.

Allman JM, Tetreault NA, Hakeem AY, Park S (2011) The von Economo Neurons in Apes and Humans. American Journal of Human Biology 23:5-21.

Allman JM, Watson KK, Tetreault NA, Hakeem AY (2005) Intuition and autism: a possible role for Von Economo neurons. Trends in Cognitive Sciences 9:367-373.

Amaral DG, Price JL (1984) Amygdalo-cortical projections in the monkey (macaca-fascicularis). Journal of Comparative Neurology 230:465-496.

Amiez C, Joseph JP, Procyk E (2005) Anterior cingulate error-related activity is modulated by predicted reward. European Journal of Neuroscience 21:3447-3452.

Amodio DM, Frith CD (2006) Meeting of minds: the medial frontal cortex and social cognition. Nature Reviews Neuroscience 7:268-277.

Anderson SW, Bechara A, Damasio H, Tranel D, Damasio AR (1999) Impairment of social and moral behavior related to early damage in human prefrontal cortex. Nature Neuroscience 2:1032-1037.

Andersson JLR, Hutton C, Ashburner J, Turner R, Friston K (2001) Modeling geometric deformations in EPI time series. Neuroimage 13:903-919.

Apicella P, Scarnati E, Schultz W (1991) Tonically discharging neurons of monkey striatum respond to preparatory and rewarding stimuli. Experimental Brain Research 84:672-675.

Apperly IA, Butterfill SA (2009) Do Humans Have Two Systems to Track Beliefs and Belief-Like States? Psychological Review 116:953-970.

Apps MAJ, Balsters JH, Ramnani N (accepted) The anterior cingulate cortex: monitoring the outcomes of others' decisions. Social Neuroscience.

Asch SE (1956) Studies of independence and conformity: A minority of one against a unanimous majority. In. Psychological Monographs.

Ashburner J, Friston KJ (2005) Unified segmentation. Neuroimage 26:839-851.

Ballantine HT, Bouckoms AJ, Thomas EK, Giriunas IE (1987) Treatment of psychiatric-illness by stereotaxic cingulotomy. Biological Psychiatry 22:807-819.

Ballantine HT, Jr., Cassidy WL, Flanagan NB, Marino R, Jr. (1967) Stereotaxic anterior cingulotomy for neuropsychiatric illness and intractable pain. Journal of neurosurgery 26:488-495.

Balsters JH, Ramnani N (2008) Symbolic representations of action in the human cerebellum. Neuroimage 43:388-398.

Bar M (2007) The proactive brain: using analogies and associations to generate predictions. Trends in Cognitive Sciences 11:280-289.

Bar M (2009) The proactive brain: memory for predictions. Philosophical Transactions of the Royal Society B-Biological Sciences 364:1235-1243.

Barbas H, Ghashghaei H, Dombrowski SM, Rempel-Clower NL (1999) Medial prefrontal cortices are unified by common connections with superior temporal cortices and distinguished by input from memory-related areas in the rhesus monkey. Journal of Comparative Neurology 410:343-367.

Barch DM, Braver TS, Sabb FW, Noll DC (2000) Anterior cingulate and the monitoring of response conflict: Evidence from an fMRI study of overt verb generation. Journal of Cognitive Neuroscience 12:298-309.

Bar-On R, Tranel D, Denburg NL, Bechara A (2003) Exploring the neurological substrate of emotional and social intelligence. Brain 126:1790-1800.

Baron-Cohen S, Leslie AM, Frith U (1985) Does the autistic-child have a theory of mind. Cognition 21:37-46.

Baron-Cohen S, Ring HA, Bullmore ET, Wheelwright S, Ashwin C, Williams SCR (2000) The amygdala theory of autism. Neuroscience and Biobehavioral Reviews 24:355-364.

Barrash J, Tranel D, Anderson SW (2000) Acquired personality disturbances associated with bilateral damage to the ventromedial prefrontal region. Developmental Neuropsychology 18:355-381.

Baumgartner T, Fischbacher U, Feierabend A, Lutz K, Fehr E (2009) The Neural Circuitry of a Broken Promise. Neuron 64:756-770.

Bautista LM, Tinbergen J, Kacelnik A (2001) To walk or to fly? How birds choose among foraging modes. Proceedings of the National Academy of Sciences of the United States of America 98:1089-1094.

Bautista LM, Tinbergen J, Wiersma P, Kacelnik A (1998) Optimal foraging and beyond: How starlings cope with changes in food availability. American Naturalist 152:543-561.

Bayer HM, Glimcher PW (2005) Midbrain dopamine neurons encode a quantitative reward prediction error signal. Neuron 47:129-141.

Becerra L, Breiter HC, Wise R, Gonzalez RG, Borsook D (2001) Reward circuitry activation by noxious thermal stimuli. Neuron 32:927-946.

Beckmann CF, DeLuca M, Devlin JT, Smith SM (2005) Investigations into resting-state connectivity using independent component analysis. Philosophical Transactions of the Royal Society B-Biological Sciences 360:1001-1013.

Beckmann CF, Smith SA (2004) Probabilistic independent component analysis for functional magnetic resonance imaging. Ieee Transactions on Medical Imaging 23:137-152.

Beckmann M, Johansen-Berg H, Rushworth MFS (2009) Connectivity-Based Parcellation of Human Cingulate Cortex and Its Relation to Functional Specialization. Journal of Neuroscience 29:1175-1190.

Behrens TEJ, Hunt LT, Woolrich MW, Rushworth MFS (2008) Associative learning of social value. Nature 456:245-U245.

Behrens TEJ, Woolrich MW, Walton ME, Rushworth MFS (2007) Learning the value of information in an uncertain world. Nature Neuroscience 10:1214-1221.

Bellebaum C, Kobza S, Thiele S, Daum I (2010) It Was Not MY Fault: Event-Related Brain Potentials in Active and Observational Learning from Feedback. Cerebral Cortex 20:2874-2883.

Bernoulli D, (1954) Exposition of a New Theory on the Measurement of Risk. Econometrica 22: 23-36.

Berns GS, Chappelow J, Zink CF, Pagnoni G, Martin-Skurski ME, Richards J (2005) Neurobiological correlates of social conformity and independence during mental rotation. Biological Psychiatry 58:245-253.

Bernstein C, Kacelnik A, Krebs JR (1991) Individual decisions and the distribution of predators in a patchy environment .2. the influence of travel costs and structure of the environment. Journal of Animal Ecology 60:205-225.

Berthoz S, Armony JL, Blair RJR, Dolan RJ (2002) An fMRI study of intentional and unintentional (embarrassing) violations of social norms. Brain 125:1696-1708.

Berthoz S, Grezes J, Armony JL, Passingham RE, Dolan RJ (2006) Affective response to one's own moral violations. Neuroimage 31:945-950.

Bickel WK, Pitcock JA, Yi R, Angtuaco EJC (2009) Congruence of BOLD Response across Intertemporal Choice Conditions: Fictive and Real Money Gains and Losses. Journal of Neuroscience 29:8839-8846.

Bond R, Smith PB (1996) Culture and conformity: A meta-analysis of studies using Asch's (1952b, 1956) Line judgment task. Psychological Bulletin 119:111-137.

Botvinick M, Nystrom LE, Fissell K, Carter CS, Cohen JD (1999) Conflict monitoring versus selection-for-action in anterior cingulate cortex. Nature 402:179-181.

Botvinick MM (2007) Conflict monitoring and decision making: Reconciling two perspectives on anterior cingulate function. Cognitive Affective & Behavioral Neuroscience 7:356-366.

Botvinick MM, Huffstetler S, McGuire JT (2009) Effort discounting in human nucleus accumbens. Cognitive Affective & Behavioral Neuroscience 9:16-27.

Bozkurt A, Zilles K, Schleicher A, Kamper L, Arigita ES, Uylings HBM, Kotter R (2005) Distributions of transmitter receptors in the macaque cingulate cortex. Neuroimage 25:219-229.

Brannan S, Liotti M, Egan G, Shade R, Madden L, Robillard R, Abplanalp B, Stofer K, Denton D, Fox PT (2001) Neuroimaging of cerebral activations and deactivations associated with hypercapnia and hunger for air. Proceedings of the National Academy of Sciences of the United States of America 98:2029-2034.

Braver TS, Barch DM, Gray JR, Molfese DL, Snyder A (2001) Anterior cingulate cortex and response conflict: Effects of frequency, inhibition and errors. Cerebral Cortex 11:825-836.

Bray S, O'Doherty J (2007) Neural coding of reward-prediction error signals during classical conditioning with attractive faces. Journal of Neurophysiology 97:3036-3045.

Brodmann K (1909) Vergleichende Lokalisationslehre der Groshirnrinde. Leipzig: Barth.

Brovelli A, Laksiri N, Nazarian B, Meunier M, Boussaoud D (2008) Understanding the neural computations of arbitrary visuomotor learning through fMRI and associative learning theory. Cerebral Cortex 18:1485-1495.

Brunet E, Sarfati Y, Hardy-Bayle MC, Decety J (2000) A PET investigation of the attribution of intentions with a nonverbal task. Neuroimage 11:157-166.

Brunner D, Kacelnik A, Gibbon J (1992) Optimal foraging and timing processes in the starling, sturnus-vulgaris - effect of inter-capture interval. Animal Behaviour 44:597-613.

Buccino G, Binkofski F, Fink GR, Fadiga L, Fogassi L, Gallese V, Seitz RJ, Zilles K, Rizzolatti G, Freund HJ (2001) Action observation activates premotor and parietal areas in a somatotopic manner: an fMRI study. European Journal of Neuroscience 13:400-404.

Buccino G, Vogt S, Ritzl A, Fink GR, Zilles K, Freund HJ, Rizzolatti G (2004) Neural circuits underlying imitation learning of hand actions: An event-related fMRI study. Neuron 42:323-334.

Buchel C, Bornhovd K, Quante M, Glauche V, Bromm B, Weiller C (2002) Dissociable neural responses related to pain intensity, stimulus intensity, and stimulus awareness within the anterior cingulate cortex: A parametric single-trial laser functional magnetic resonance imaging study. Journal of Neuroscience 22:970-976.

Buckholtz JW, Asplund CL, Dux PE, Zald DH, Gore JC, Jones OD, Marois R (2008) The Neural Correlates of Third-Party Punishment. Neuron 60:930-940.

Buckwalter JA, Van Hoesen GW, Parvizi J (2002) Afferent connections of the primate posterior cingulate, retrosplenial, and mesial parietal cortices. Society for Neuroscience Abstract.

Buge A, Escourolle R, Rancurel G, Poisson M (1975) Akinetic-mutism and bilateral softening of cingulate gyrus - 3 pathological cases. Revue Neurologique 131:121-137.

Burke CJ, Tobler PN, Baddeley M, Schultz W (2010) Neural mechanisms of observational learning. Proc Natl Acad Sci U S A 107:14431-14436.

Bush G, Luu P, Posner MI (2000) Cognitive and emotional influences in anterior cingulate cortex. Trends in Cognitive Sciences 4:215-222.

Bush G, Whalen PJ, Shin LM, Rauch SL (2006) The counting Stroop: a cognitive interference task. Nature Protocols 1:230-233.

Butti C, Sherwood CC, Hakeem AV, Allman JM, Hof PR (2009) Total Number and Volume of Von Economo Neurons in the Cerebral Cortex of Cetaceans. Journal of Comparative Neurology 515:243-259.

Cardinal RN (2006) Neural systems implicated in delayed and probabilistic reinforcement. Neural Networks 19:1277-1301.

Carmichael ST, Price JL (1995) Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. Journal of Comparative Neurology 363:615-641.

Caro TM, Hauser MD (1992) Is there teaching in nonhuman animals? Q Rev Biol 67:151-174.

Carter CS, Braver TS, Barch DM, Botvinick MM, Noll D, Cohen JD (1998) Anterior cingulate cortex, error detection, and the online monitoring of performance. Science 280:747-749.

Castelli F, Happe F, Frith U, Frith C (2000) Movement and mind: A functional imaging study of perception and interpretation of complex intentional movement patterns. Neuroimage 12:314-325.

Cauda F, Geminiani G, D'Agata F, Sacco K, Duca S, Bagshaw AP, Cavanna AE (2010) Functional Connectivity of the Posteromedial Cortex. Plos One 5.

Cavada C, Company T, Tejedor J, Cruz-Rizzolo RJ, Reinoso-Suarez F (2000) The anatomical connections of the macaque monkey orbitofrontal cortex. A review. Cerebral Cortex 10:220-242.

Chaminade T, Hodgins J, Kawato M (2007) Anthropomorphism influences perception of computer-animated characters' actions. Social Cognitive and Affective Neuroscience 2:206-216.

Charnov EL (1976) Optimal foraging, marginal value theorem. Theoretical Population Biology 9:129-136.

Chib VS, Rangel A, Shimojo S, O'Doherty JP (2009) Evidence for a Common Representation of Decision Values for Dissimilar Goods in Human Ventromedial Prefrontal Cortex. Journal of Neuroscience 29:12315-12320.

Cialdini RB, Goldstein NJ (2004) Social influence: Compliance and conformity. Annual Review of Psychology 55:591-621.

Cloutier J, Gabrieli JDE, O'Young D, Ambady N (2011) An fMRI study of violations of social expectations: When people are not who we expect them to be. Neuroimage 57:583-588.

Cohen RA, Kaplan RF, Moser DJ, Jenkins MA, Wilkinson H (1999) Impairments of attention after cingulotomy. Neurology 53:819-824.

Cohen RA, Paul R, Zawacki TM, Moser DJ, Sweet L, Wilkinson H (2001) Emotional and personality changes following cingulotomy. Emotion 1:38-50.

Cole MW, Yeung N, Freiwald WA, Botvinick M (2009) Cingulate cortex: Diverging data from humans and monkeys. Trends in Neurosciences 32:566-574.

Cole MW, Yeung N, Freiwald WA, Botvinick M (2010) Conflict over Cingulate Cortex: Between-Species Differences in Cingulate May Support Enhanced Cognitive Flexibility in Humans. Brain Behavior and Evolution 75:239-240.

Corbetta M, Miezin FM, Dobmeyer S, Shulman GL, Petersen SE (1991) Selective and divided attention during visual discriminations of shape, color, and speed - functional-anatomy by positron emission tomography. Journal of Neuroscience 11:2383-2402.

Croxson PL, Walton ME, O'Reilly JX, Behrens TEJ, Rushworth MFS (2009) Effort-Based Cost-Benefit Valuation and the Human Brain. Journal of Neuroscience 29:4531-4541.

Cusack R, Brett M, Osswald K (2003) An evaluation of the use of magnetic field maps to undistort echo-planar images. Neuroimage 18:127-142.

Dalley JW, Everitt BJ, Robbins TW (2011) Impulsivity, Compulsivity, and Top-Down Cognitive Control. Neuron 69:680-694.

Damasio AR, Geschwind N (1984) the neural basis of language. Annual Review of Neuroscience 7:127-147.

Damasio AR, Van Hoesen GW (1983) Focal lesions of the limbic frontal lobe. In: Neuropsychology of human emotion (Heilman KM, Satz P, eds), pp 85-110. New York: Guilford Press.

D'Ardenne K, McClure SM, Nystrom LE, Cohen JD (2008) BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. Science 319:1264-1267.

Davis KD, Taylor KS, Hutchison WD, Dostrovsky JO, McAndrews MP, Richter EO, Lozano AM (2005) Human anterior cingulate cortex neurons encode cognitive and emotional demands. Journal of Neuroscience 25:8402-8406.

Day JJ, Jones JL, Wightman RM, Carelli RM (2010) Phasic Nucleus Accumbens Dopamine Release Encodes Effort- and Delay-Related Costs. Biological Psychiatry 68:306-309.

Dayan P, Balleine BW (2002) Reward, motivation, and reinforcement learning. Neuron 36:285-298.

Dayan P, Daw ND (2008) Decision theory, reinforcement learning, and the brain. Cognitive Affective & Behavioral Neuroscience 8:429-453.

de Bruijn ERA, de Lange FP, von Cramon DY, Ullsperger M (2009) When Errors Are Rewarding. Journal of Neuroscience 29:12183-12186.

Deichmann R, Gottfried JA, Hutton C, Turner R (2003) Optimized EPI for fMRI studies of the orbitofrontal cortex. Neuroimage 19:430-441.

Deichmann R, Josephs O, Hutton C, Corfield DR, Turner R (2002) Compensation of susceptibility-induced BOLD sensitivity losses in echo-planar fMRI Imaging. Neuroimage 15:120-135.

den Ouden HEM, Friston KJ, Daw ND, McIntosh AR, Stephan KE (2009) A Dual Role for Prediction Error in Associative Learning. Cerebral Cortex 19:1175-1185.

Denk F, Walton ME, Jennings KA, Sharp T, Rushworth MFS, Bannerman DM (2005) Differential involvement of serotonin and dopamine systems in cost-benefit decisions about delay or effort. Psychopharmacology 179:587-596.

Devinsky O, Morrell MJ, Vogt BA (1995) Contributions of anterior cingulate cortex to behaviour. Brain 118:279-306.

Doya K (2008) Modulators of decision making. Nature Neuroscience 11:410-416.

Dum RP, Strick PL (1991) The origin of corticospinal projections from the premotor areas in the frontal-lobe. Journal of Neuroscience 11:667-689.

Dum RP, Strick PL (1996) Spinal cord terminations of the medial wall motor areas in macaque monkeys. Journal of Neuroscience 16:6513-6525.

Dunbar RIM, Shultz S (2007) Evolution in the social brain. Science 317:1344-1347.

Dunnett CW (1970) Multiple comparison tests. Biometrics 26:139.

Fan J, Hof PR, Guise KG, Fossella JA, Posner MI (2008) The functional integration of the anterior cingulate cortex during conflict processing. Cerebral Cortex 18:796-805.

Fehr E, Fischbacher U (2004) Social norms and human cooperation. Trends in Cognitive Sciences 8:185-190.

Fiorillo CD, Tobler PN, Schultz W (2003) Discrete coding of reward probability and uncertainty by dopamine neurons. Science 299:1898-1902.

Fletcher PC, Happe F, Frith U, Baker SC, Dolan RJ, Frackowiak RSJ, Frith CD (1995) Other minds in the brain - a functional imaging study of theory of mind in story comprehension. Cognition 57:109-128.

Floresco SB, Ghods-Sharifi S (2007) Amygdala-prefrontal cortical circuitry regulates effort-based decision making. Cerebral Cortex 17:251-260.

Floresco SB, Tse MTL, Ghods-Sharifi S (2008) Dopaminergic and glutamatergic regulation of effort- and delay-based decision making. Neuropsychopharmacology 33:1966-1979.

Fogassi L, Ferrari PF, Gesierich B, Rozzi S, Chersi F, Rizzolatti G (2005) Parietal lobe: From action organization to intention understanding. Science 308:662-667.

Frank MJ, Woroch BS, Curran T (2005) Error-related negativity predicts reinforcement learning and conflict biases. Neuron 47:495-501.

Franks NR, Richardson T (2006) Teaching in tandem-running ants. Nature 439:153.

Freeman KB, Green L, Myerson J, Woolverton WL (2009) Delay discounting of saccharin in rhesus monkeys. Behavioural Processes 82:214-218.

Friederici AD (2002) Towards a neural basis of auditory sentence processing. Trends in Cognitive Sciences 6:78-84.

Friston KJ, Ashburner J, Frith CD, Poline JB, Heather JD, Frackowiak RSJ (1995a) Spatial registration and normalization of images. Human Brain Mapping 3:165-189.

Friston KJ, Holmes AP, Poline JB, Grasby PJ, Williams SCR, Frackowiak RSJ, Turner R (1995b) Analysis of fmri time-series revisited. Neuroimage 2:45-53.

Friston KJ, Stephan KE, Lund TE, Morcom A, Kiebel S (2005) Mixed-effects and fMRI studies. Neuroimage 24:244-252.

Friston KJ, Williams S, Howard R, Frackowiak RSJ, Turner R (1996) Movement-related effects in fMRI time-series. Magnetic Resonance in Medicine 35:346-355.

Frith CD, Frith U (1999) Cognitive psychology - Interacting minds - A biological basis. Science 286:1692-1695.

Frith CD, Frith U (2006) The neural basis of mentalizing. Neuron 50:531-534.

Frith U, Frith C (2010) The social brain: allowing humans to boldly go where no other species has been. Philosophical Transactions of the Royal Society B-Biological Sciences 365:165-175.

Frith U, Frith CD (2003) Development and neurophysiology of mentalizing. Philosophical Transactions of the Royal Society B-Biological Sciences 358:459-473.

Frysztak RJ, Neafsey EJ (1994) The effect of medial frontal-cortex lesions on cardiovascular conditioned emotional responses in the rat. Brain Research 643:181-193.

Gallagher HL, Happe F, Brunswick N, Fletcher PC, Frith U, Frith CD (2000) Reading the mind in cartoons and stories: an fMRI study of 'theory of mind' in verbal and nonverbal tasks. Neuropsychologia 38:11-21.

Gallese V (2007) Before and below 'theory of mind': embodied simulation and the neural correlates of social cognition. Philosophical Transactions of the Royal Society B-Biological Sciences 362:659-669.

Gallese V, Fadiga L, Fogassi L, Rizzolatti G (1996) Action recognition in the premotor cortex. Brain 119:593-609.

Gallese V, Goldman A (1998) Mirror neurons and the simulation theory of mind-reading. Trends in Cognitive Sciences 2:493-501.

Garavan H, Ross TJ, Murphy K, Roche RAP, Stein EA (2002) Dissociable executive functions in the dynamic control of behavior: Inhibition, error detection, and correction. Neuroimage 17:1820-1829.

Gazzola V, Keysers C (2009) The Observation and Execution of Actions Share Motor and Somatosensory Voxels in all Tested Subjects: Single-Subject Analyses of Unsmoothed fMRI Data. Cerebral Cortex 19:1239-1255.

Genovese CR, Lazar NA, Nichols T (2002) Thresholding of statistical maps in functional neuroimaging using the false discovery rate. Neuroimage 15:870-878.

Glascher J, Daw N, Dayan P, O'Doherty JP (2010) States versus Rewards: Dissociable Neural Prediction Error Signals Underlying Model-Based and Model-Free Reinforcement Learning. Neuron 66:585-595.

Glover GH (1999) 3D z-shim method for reduction of susceptibility effects in BOLD fMRI. Magnetic Resonance in Medicine 42:290-299.

Green L, Fristoe N, Myerson J (1994a) Temporal discounting and preference reversals in choice between delayed outcomes. Psychonomic Bulletin & Review 1:383-389.

Green L, Fry AF, Myerson J (1994b) Discounting of delayed rewards - a life-span comparison. Psychological Science 5:33-36.

Green L, Myerson J, McFadden E (1997) Rate of temporal discounting decreases with amount of reward. Memory & Cognition 25:715-723.

Grezes J, Frith CD, Passingham RE (2004) Inferring false beliefs from the actions of oneself and others: an fMRI study. Neuroimage 21:744-750.

Groenewegen HJ, Berendse HW, Haber SN (1993) Organization of the output of the ventral striatopallidal system in the rat - ventral pallidal efferents. Neuroscience 57:113-142.

Grootoonk S, Hutton C, Ashburner J, Howseman AM, Josephs O, Rees G, Friston KJ, Turner R (2000) Characterization and correction of interpolation effects in the realignment of fMRI time series. Neuroimage 11:49-57.

Haber SN, Fudge JL (1997) The primate substantia nigra and VTA: Integrative circuitry and function. Critical Reviews in Neurobiology 11:323-342.

Haber SN, Knutson B (2010) The Reward Circuit: Linking Primate Anatomy and Human Imaging. Neuropsychopharmacology 35:4-26.

Haber SN, Kunishio K, Mizobuchi M, Lyndbalta E (1995) The orbital and medial prefrontal circuit through the primate basal ganglia. Journal of Neuroscience 15:4851-4867.

Hadland KA, Rushworth MFS, Gaffan D, Passingham RE (2003) The effect of cingulate lesions on social behaviour and emotion. Neuropsychologia 41:919-931.

Hampton AN, Bossaerts P, O'Doherty JP (2006) The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. Journal of Neuroscience 26:8360-8367.

Hampton AN, Bossaerts P, O'Doherty JP (2008) Neural correlates of mentalizing-related computations during strategic interactions in humans. Proceedings of the National Academy of Sciences of the United States of America 105:6741-6746.

Hampton AN, O'Doherty JP (2007) Decoding the neural substrates of reward-related decision making with functional MRI. Proceedings of the National Academy of Sciences of the United States of America 104:1377-1382.

Handwerker DA, Ollinger JM, D'Esposito M (2004) Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. Neuroimage 21:1639-1651.

Hare TA, Camerer CF, Rangel A (2009) Self-Control in Decision-Making Involves Modulation of the vmPFC Valuation System. Science 324:646-648.

Hare TA, O'Doherty J, Camerer CF, Schultz W, Rangel A (2008) Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. Journal of Neuroscience 28:5623-5630.

Hauber W, Sommer S (2009) Prefrontostriatal Circuitry Regulates Effort-Related Decision Making. Cerebral Cortex 19:2240-2247.

Hayden BY, Heilbronner SR, Pearson JM, Platt ML (2011b) Surprise Signals in Anterior Cingulate Cortex: Neuronal Encoding of Unsigned Reward Prediction Errors Driving Adjustment in Behavior. Journal of Neuroscience 31:4178-4187.

Hayden BY, Nair AC, McCoy AN, Platt ML (2008) Posterior Cingulate Cortex Mediates Outcome-Contingent Allocation of Behavior. Neuron 60:19-25.

Hayden BY, Pearson JM, Platt ML (2011a) Neuronal basis of sequential foraging decisions in a patchy environment. Nature Neuroscience 14:933-U165.

Hayden BY, Platt ML (2010) Neurons in Anterior Cingulate Cortex Multiplex Information about Reward and Action. Journal of Neuroscience 30:3339-3346.

He SQ, Dum RP, Strick PL (1995) Topographic organization of corticospinal projections from the frontal-lobe - motor areas on the medial surface of the hemisphere. Journal of Neuroscience 15:3284-3306.

Heilbronner SR, Hayden BY, Platt ML (2011) Decision salience signals in posterior cingulate cortex. Frontiers in neuroscience 5:55.

Henson R, Rugg MD (2001) Effects of stimulus repetition on latency of BOLD impulse response. Neuroimage 13:S683-S683.

Henson RNA, Price CJ, Rugg MD, Turner R, Friston KJ (2002) Detecting latency differences in event-related BOLD responses: Application to words versus nonwords and initial versus repeated face presentations. Neuroimage 15:83-97.

Hillman KL, Bilkey DK (2010) Neurons in the Rat Anterior Cingulate Cortex Dynamically Encode Cost-Benefit in a Spatial Decision-Making Task. Journal of Neuroscience 30:7705-7713.

Holekamp KE, Sakai ST, Lundrigan BL (2007) The spotted hyena (Crocuta crocuta) as a model system for study of the evolution of intelligence. Journal of Mammalogy 88:545-554.

Hollerman JR, Schultz W (1998) Dopamine neurons report an error in the temporal prediction of reward during learning. Nature Neuroscience 1:304-309.

Holroyd CB, Coles MGH (2002) The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. Psychological Review 109:679-709.

Holroyd CB, Coles MGH (2008) Dorsal anterior cingulate cortex integrates reinforcement history to guide voluntary behaviour. Cortex 44:548-559.

Holroyd CB, Krigolson OE, Baker R, Lee S, Gibson J (2009) When is an error not a prediction error? An electrophysiological investigation. Cognitive Affective & Behavioral Neuroscience 9:59-70.

Holroyd CB, Nieuwenhuis S, Yeung N, Cohen JD (2003) Errors in reward prediction are reflected in the event-related brain potential. Neuroreport 14:2481-2484.

Holroyd CB, Nieuwenhuis S, Yeung N, Nystrom L, Mars RB, Coles MGH, Cohen JD (2004) Dorsal anterior cingulate cortex shows fMRI response to internal and external error signals. Nature Neuroscience 7:497-498.

Holroyd CB, Yeung N, Coles MGH, Cohen JD (2005) A mechanism for error detection in speeded response time tasks. Journal of Experimental Psychology-General 134:163-191.

Hooker CI, Verosky SC, Germine LT, Knight RT, D'Esposito M (2008) Mentalizing about emotion and its relationship to empathy. Social Cognitive and Affective Neuroscience 3:204-217.

Hoppitt WJE, Brown GR, Kendal R, Rendell L, Thornton A, Webster MM, Laland KN (2008) Lessons from animal teaching. Trends Ecol Evol 23:486-493.

Hsu JJ, Glover GH (2005) Mitigation of susceptibility-induced signal loss in neuroimaging using localized shim coils. Magnetic Resonance in Medicine 53:243-248.

Hughes C, Adlam A, Happe F, Jackson J, Taylor A, Caspi A (2000) Good test-retest reliability for standard and advanced false-belief tasks across a wide range of abilities. Journal of Child Psychology and Psychiatry and Allied Disciplines 41:483-490.

Hughes G, Yeung N (2011) Dissociable correlates of response conflict and error awareness in error-related brain activity. Neuropsychologia 49:405-415.

Hurleygius KM, Neafsey EJ (1986) The medial frontal-cortex and gastric-motility - microstimulation results and their possible significance for the overall pattern of organization of rat frontal and parietal cortex. Brain Research 365:241-248.

Hutchins KD, Martino AM, Strick PL (1988) Corticospinal projections from the medial wall of the hemisphere. Experimental Brain Research 71:667-672.

Hutchison WD, Davis KD, Lozano AM, Tasker RR, Dostrovsky JO (1999) Pain-related neurons in the human cingulate cortex. Nature Neuroscience 2:403-405.

Hutton C, Bork A, Josephs O, Deichmann R, Ashburner J, Turner R (2002) Image distortion correction in fMRI: A quantitative evaluation. Neuroimage 16:217-240.

Hwang J, Kim S, Lee D (2009) Temporal discounting and inter-temporal choice in rhesus monkeys. Frontiers in behavioral neuroscience 3:9.

Iacoboni M, Dapretto M (2006) The mirror neuron system and the consequences of its dysfunction. Nature Reviews Neuroscience 7:942-951.

Iacoboni M, Molnar-Szakacs I, Gallese V, Buccino G, Mazziotta JC, Rizzolatti G (2005) Grasping the intentions of others with one's own mirror neuron system. Plos Biology 3:529-535.

Ishiwari K, Weber SM, Mingote S, Correa M, Salamone JD (2004) Accumbens dopamine and the regulation of effort in food-seeking behavior: modulation of work output by different ratio or force requirements. Behavioural Brain Research 151:83-91.

Iwata K, Kamo H, Ogawa A, Tsuboi Y, Noma N, Mitsuhashi Y, Taira M, Koshikawa N, Kitagawa J (2005) Anterior Cingulate cortical neuronal activity during perception of noxious thermal stimuli in monkeys. Journal of Neurophysiology 94:1980-1991.

Jabbi M, Bastiaansen J, Keysers C (2008) A Common Anterior Insula Representation of Disgust Observation, Experience and Imagination Shows Divergent Functional Connectivity Pathways. Plos One 3:8.

Jackson PL, Meltzoff AN, Decety J (2006) Neural circuits involved in imitation and perspective-taking. Neuroimage 31:429-439.

Janer KW, Pardo JV (1991) Deficits in selective attention following bilateral anterior cingulotomy. Journal of Cognitive Neuroscience 3:231-241.

Jenike MA, Baer L, Ballantine T, Martuza RL, Tynes S, Giriunas I, Buttolph ML, Cassem NH (1991) Cingulotomy for refractory obsessive-compulsive disorder - a long-term follow-up of 33 patients. Archives of General Psychiatry 48:548-555.

Jessup RK, Busemeyer JR, Brown JW (2010) Error Effects in Anterior Cingulate Cortex Reverse when Error Likelihood Is High. Journal of Neuroscience 30:3467-3472.

Jimura K, Myerson J, Hilgard J, Braver TS, Green L (2009) Are people really more patient than other animals? Evidence from human discounting of real liquid rewards. Psychonomic Bulletin & Review 16:1071-1075.

Jocham G, Neumann J, Klein TA, Danielmeier C, Ullsperger M (2009) Adaptive Coding of Action Values in the Human Rostral Cingulate Zone. Journal of Neuroscience 29:7489-7496.

Johnson MW, Bickel WK (2002) Within-subject comparison of real and hypothetical money rewards in delay discounting. Journal of the Experimental Analysis of Behavior 77:129-146.

Kaada BR, Pribram KH, Epstein JA (1949) Respiratory and vascular responses in monkeys from temporal pole, insula, orbital surface and cingulate gyrus; a preliminary report. Journal of neurophysiology 12:347-356.

Kable JW, Glimcher PW (2007) The neural correlates of subjective value during intertemporal choice. Nature Neuroscience 10:1625-1633.

Kable JW, Glimcher PW (2010) An "As Soon As Possible" Effect in Human Intertemporal Decision Making: Behavioral Evidence and Neural Mechanisms. Journal of Neurophysiology 103:2513-2531.

Kang SK, Hirsh JB, Chasteen AL (2010) Your mistakes are mine: Self-other overlap predicts neural response to observed errors. Journal of Experimental Social Psychology 46:229-232.

Kawashima R, Inoue K, Sugiura M, Okada K, Ogawa A, Fukuda H (1999) A positron emission tomography study of self-paced finger movements at different frequencies. Neuroscience 92:107-112.

Kawashima R, Satoh K, Itoh H, Ono S, Furumoto S, Gotoh R, Koyama M, Yoshioka S, Takahashi T, Takahashi K, Yanagisawa T, Fukuda H (1996) Functional anatomy of GO/NO-GO discrimination and response selection - A PET study in man. Brain Research 728:79-89.

Kawato M, Samejima K (2007) Efficient reinforcement learning: computational theories, neuroscience and robotics. Current Opinion in Neurobiology 17:205-212.

Kennerley SW, Dahmubed AF, Lara AH, Wallis JD (2009) Neurons in the Frontal Lobe Encode the Value of Multiple Decision Variables. Journal of Cognitive Neuroscience 21:1162-1178.

Kennerley SW, Wallis JD (2009a) Encoding of Reward and Space During a Working Memory Task in the Orbitofrontal Cortex and Anterior Cingulate Sulcus. Journal of Neurophysiology 102:3352-3364.

Kennerley SW, Wallis JD (2009b) Evaluating choices by single neurons in the frontal lobe: outcome value encoded across multiple decision variables. European Journal of Neuroscience 29:2061-2073.

Kennerley SW, Walton ME, Behrens TEJ, Buckley MJ, Rushworth MFS (2006) Optimal decision making and the anterior cingulate cortex. Nature Neuroscience 9:940-947.

Kerns JG, Cohen JD, MacDonald AW, Cho RY, Stenger VA, Carter CS (2004) Anterior Cingulate conflict monitoring and adjustments in control. Science 303:1023-1026.

Keysers C, Gazzola V (2007) Integrating simulation and theory of mind: from self to social cognition. Trends in Cognitive Sciences 11:194-196.

Kiehl KA, Liddle PF, Hopfinger JB (2000) Error processing and the rostral anterior cingulate: An event-related fMRI study. Psychophysiology 37:216-223.

Kilner JM (2011) More than one pathway to action understanding. Trends in cognitive sciences 15:352-357.

Kilner JM, Neal A, Weiskopf N, Friston KJ, Frith CD (2009) Evidence of Mirror Neurons in Human Inferior Frontal Gyrus. Journal of Neuroscience 29:10153-10159.

Kim S, Hwang J, Lee D (2008) Prefrontal coding of temporally discounted values during intertemporal choice. Neuron 59:161-172.

King-Casas B, Tomlin D, Anen C, Camerer CF, Quartz SR, Montague PR (2005) Getting to know you: Reputation and trust in a two-person economic exchange. Science 308:78-83.

Kirby KN (1997) Bidding on the future: Evidence against normative discounting of delayed rewards. Journal of Experimental Psychology-General 126:54-70.

Kirby KN, Marakovic NN (1996) Delay-discounting probabilistic rewards: Rates decrease as amounts increase. Psychonomic Bulletin & Review 3:100-104.

Klein A, Andersson J, Ardekani BA, Ashburner J, Avants B, Chiang MC, Christensen GE, Collins DL, Gee J, Hellier P, Song JH, Jenkinson M, Lepage C, Rueckert D, Thompson P, Vercauteren T, Woods RP, Mann JJ, Parsey RV (2009) Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. Neuroimage 46:786-802.

Klucharev V, Hytonen K, Rijpkema M, Smidts A, Fernandez G (2009) Reinforcement Learning Signal Predicts Social Conformity. Neuron 61:140-151.

Knutson B, Bossaerts P (2007) Neural antecedents of financial decisions. Journal of Neuroscience 27:8174-8177.

Knutson B, Cooper JC (2005) Functional magnetic resonance imaging of reward prediction. Current Opinion in Neurology 18:411-417.

Knutson B, Westdorp A, Kaiser E, Hommer D (2000) FMRI visualization of brain activity during a monetary incentive delay task. Neuroimage 12:20-27.

Koban L, Pourtois G, Vocat R, Vuilleumier P (2010) When your errors make me lose or win: Event-related potentials to observed errors of cooperators and competitors. Social Neuroscience 5:360-374.

Kobayashi S, Schultz W (2008) Influence of reward delays on responses of dopamine neurons. Journal of Neuroscience 28:7837-7846.

Koski L, Paus T (2000) Functional connectivity of the anterior cingulate cortex within the human frontal lobe: a brain-mapping meta-analysis (vol 133, pg 55, 2000). Experimental Brain Research 134:408-408.

Kunishio K, Haber SN (1994) Primate cingulostriatal projection - limbic striatal versus sensorimotor striatal input. Journal of Comparative Neurology 350:337-356.

Kurniawan IT, Seymour B, Talmi D, Yoshida W, Chater N, Dolan RJ (2010) Choosing to Make an Effort: The Role of Striatum in Signaling Physical Effort of a Chosen Action. Journal of Neurophysiology 104:313-321.

Lau HC, Rogers RD, Haggard P, Passingham RE (2004a) Attention to intention. Science 303:1208-1210.

Lau HC, Rogers RD, Passingham RE (2006) On measuring the perceived onsets of spontaneous actions. Journal of Neuroscience 26:7265-7271.

Lau HC, Rogers RD, Ramnani N, Passingham RE (2004b) Willed action and attention to the selection of action. Neuroimage 21:1407-1415.

Le Bihan D (2003) Looking into the functional architecture of the brain with diffusion MRI. Nature Reviews Neuroscience 4:469-480.

Le Bihan D, Mangin JF, Poupon C, Clark CA, Pappata S, Molko N, Chabriat H (2001) Diffusion tensor imaging: Concepts and applications. Journal of Magnetic Resonance Imaging 13:534-546.

Leichnetz GR (2001) Connections of the medial posterior parietal cortex (area 7m) in the monkey. Anatomical Record 263:215-236.

Lenz FA, Rios M, Zirh A, Chau D, Krauss G, Lesser RP (1998) Painful stimuli evoke potentials recorded over the human anterior cingulate gyrus. Journal of Neurophysiology 79:2231-2234.

Leslie KR, Johnson-Frey SH, Grafton ST (2004) Functional imaging of face and hand imitation: towards a motor theory of empathy. Neuroimage 21:601-607.

Liotti M, Brannan S, Egan G, Shade R, Madden L, Abplanalp B, Robillard R, Lancaster J, Zamarripa FE, Fox PT, Denton D (2001) Brain responses associated with consciousness of breathlessness (air hunger). Proceedings of the National Academy of Sciences of the United States of America 98:2035-2040.

Lipschutz B, Friston KJ, Ashburner J, Turner R, Price CJ (2001) Assessing study-specific regional variations in fMRI signal. Neuroimage 13:392-398.

Litt A, Plassmann H, Shiv B, Rangel A (2011) Dissociating Valuation and Saliency Signals during Decision-Making. Cerebral Cortex 21:95-102.

Liu D, Sabbagh MA, Gehring WJ, Wellman HM (2009) Neural Correlates of Children's Theory of Mind Development. Child Development 80:318-326.

Louie K, Glimcher PW (2010) Separating Value from Choice: Delay Discounting Activity in the Lateral Intraparietal Area. Journal of Neuroscience 30:5498-5507.

Luhmann CC, Chun MM, Yi DJ, Lee D, Wang XJ (2008) Neural Dissociation of Delay and Uncertainty in Intertemporal Choice. Journal of Neuroscience 28:14459-14466.

Luppino G, Matelli M, Camarda RM, Gallese V, Rizzolatti G (1991) Multiple representations of body movements in mesial area-6 and the adjacent cingulate cortex - an intracortical microstimulation study in the macaque monkey. Journal of Comparative Neurology 311:463-482.

Lyndbalta E, Haber SN (1994) Primate striatonigral projections - a comparison of the sensorimotor-related striatum and the ventral striatum. Journal of Comparative Neurology 345:562-578.

Madden GJ, Begotka AM, Raiff BR, Kastern LL (2003) Delay discounting of real and hypothetical rewards. Experimental and Clinical Psychopharmacology 11:139-145.

Maier ME, Yeung N, Steinhauser M (2011) Error-related brain activity and adjustments of selective attention following errors. Neuroimage 56:2339-2347.

Malfait N, Valyear KF, Culham JC, Anton JL, Brown LE, Gribble PL (2010) fMRI Activation during Observation of Others' Reach Errors. Journal of Cognitive Neuroscience 22:1493-1503.

Mansfield P (1984) Real-time echo-planar imaging by nmr. British Medical Bulletin 40:187.

Mar AC, Walker ALJ, Theobald DE, Eagle DM, Robbins TW (2011) Dissociable Effects of Lesions to Orbitofrontal Cortex Subregions on Impulsive Choice in the Rat. Journal of Neuroscience 31:6398-6404.

Margulies DS, Kelly AMC, Uddin LQ, Biswal BB, Castellanos FX, Milham MP (2007) Mapping the functional connectivity of anterior cingulate cortex. Neuroimage 37:579-588.

Markowitsch HJ, Emmans D, Irle E, Streicher M, Preilowski B (1985) Cortical and subcortical afferent connections of the primates temporal pole - a study of rhesus-monkeys, squirrel-monkeys, and marmosets. Journal of Comparative Neurology 242:425-458.

Mars RB, Shea NJ, Kolling N, Rushworth M (in press) Model-based analyses: Promises, pitfalls, and example applications to the study of cognitive control. The Quarterly Journal of Experimental Psychology.

Matsumoto M, Matsumoto K, Abe H, Tanaka K (2007) Medial prefrontal cell activity signaling prediction errors of action values. Nature Neuroscience 10:647-656.

Mazoyer BM, Tzourio N, Frak V, Syrota A, Murayama N, Levrier O, Salamon G, Dehaene S, Cohen L, Mehler J (1993) the cortical representation of speech. Journal of Cognitive Neuroscience 5:467-479.

Mazur JE (1985) probability and delay of reinforcement as factors in discrete-trial choice. Journal of the Experimental Analysis of Behavior 43:341-351.

Mazur JE (1986) choice between single and multiple delayed reinforcers. Journal of the Experimental Analysis of Behavior 46:67-77.

Mazur JE (1997) Choice, delay, probability, and conditioned reinforcement. Animal Learning & Behavior 25:131-147.

Mazur JE (2001) Hyperbolic value addition and general models of animal choice. Psychological Review 108:96-112.

Mazur JE, Snyderman M, Coe D (1985) influences of delay and rate of reinforcement on discrete-trial choice. Journal of Experimental Psychology-Animal Behavior Processes 11:565-575.

McClure SM, Berns GS, Montague PR (2003) Temporal prediction errors in a passive learning task activate human striatum. Neuron 38:339-346.

McClure SM, Laibson DI, Loewenstein G, Cohen JD (2004) Separate neural systems value immediate and delayed monetary rewards. Science 306:503-507.

Menon V, Adleman NE, White CD, Glover GH, Reiss AL (2001) Error-related brain activation during a Go/NoGo response inhibition task. Human Brain Mapping 12:131-143.

Merboldt KD, Fransson P, Bruhn H, Frahm J (2001) Functional MRI of the human amygdala? Neuroimage 14:253-257.

Mesulam MM, Mufson EJ (1982) insula of the old-world monkey .3. Efferent cortical output and comments on function. Journal of Comparative Neurology 212:38-52.

Middleton FA, Strick PL (2000) Basal ganglia output and cognition: Evidence from anatomical, behavioral, and clinical studies. Brain and Cognition 42:183-200.

Miller SA (2009) Children's Understanding of Second-Order Mental States. Psychological Bulletin 135:749-773.

Minamimoto T, La Camera G, Richmond BJ (2009) Measuring and Modeling the Interaction Among Reward Size, Delay to Reward, and Satiation Level on Motivation in Monkeys. Journal of Neurophysiology 101:437-447.

Mingote S, Font L, Farrar AM, Vontell R, Worden LT, Stopper CM, Port RG, Sink KS, Bunce JG, Chrobak JJ, Salamone JD (2008) Nucleus accumbens adenosine A(2A) receptors regulate exertion of effort by acting on the ventral striatopallidal pathway. Journal of Neuroscience 28:9037-9046.

Mirenowicz J, Schultz W (1996) Preferential activation of midbrain dopamine neurons by appetitive rather than aversive stimuli. Nature 379:449-451.

Mobbs D, Petrovic P, Marchant JL, Hassabis D, Weiskopf N, Seymour B, Dolan RJ, Frith CD (2007) When fear is near: Threat imminence elicits prefrontal-periaqueductal gray shifts in humans. Science 317:1079-1083.

Mobbs D, Yu R, Rowe JB, Eich H, FeldmanHall O, Dalgleish T (2010) Neural activity associated with monitoring the oscillating threat value of a tarantula. Proceedings of the National Academy of Sciences of the United States of America 107:20582-20586.

Monfardini E, Brovelli A, Boussaoud D, Takerkart S, Wicker B (2008) I learned from what you did: Retrieving visuornotor associations learned by observation. Neuroimage 42:1207-1213.

Montague PR, Dayan P, Sejnowski TJ (1996) A framework for mesencephalic dopamine systems based on predictive Hebbian learning. Journal of Neuroscience 16:1936-1947.

Morecraft RJ, Geula C, Mesulam MM (1992) Cytoarchitecture and neural afferents of orbitofrontal cortex in the brain of the monkey. Journal of Comparative Neurology 323:341-358.

Morecraft RJ, Van Hoesen GW (1998) Convergence of limbic input to the cingulate motor cortex in the rhesus monkey. Brain Research Bulletin 45:209-232.

Mufson EJ, Mesulam MM (1982) Insula of the old-world monkey 2. afferent cortical input and comments on the claustrum. Journal of Comparative Neurology 212:23-37.

Myerson J, Green L, Warusawitharana M (2001) Area under the curve as a measure of discounting. Journal of the Experimental Analysis of Behavior 76:235-243.

Nakano K, Kayahara T, Tsutsumi T, Ushiro H (2000) Neural circuits and functional organization of the striatum. Journal of Neurology 247:V/1-V/15.

Neafsey EJ, Hurleygius KM, Arvanitis D (1986) The topographical organization of neurons in the rat medial frontal, insular and olfactory cortex projecting to the solitary nucleus, olfactory-bulb, periaqueductal gray and superior colliculus. Brain Research 377:261-270.

Nee DE, Kastner S, Brown JW (2011) Functional heterogeneity of conflict, error, task-switching, and unexpectedness effects within medial prefrontal cortex. Neuroimage 54:528-540.

Nee DE, Wager TD, Jonides J (2007) Interference resolution: Insights from a meta-analysis of neuroimaging tasks. Cognitive Affective & Behavioral Neuroscience 7:1-17.

Nelissen K, Luppino G, Vanduffel W, Rizzolatti G, Orban GA (2005) Observing others: Multiple action representation in the frontal lobe. Science 310:332-336.

Nemeth G, Hegedus K, Molnar L (1988) Akinetic mutism associated with bicingular lesions - clinicopathological and functional anatomical correlates. European Archives of Psychiatry and Clinical Neuroscience 237:218-222.

Nieuwenhuis S, Ridderinkhof KR, Blow J, Band GPH, Kok A (2001) Error-related brain potentials are differentially related to awareness of response errors: Evidence from an antisaccade task. Psychophysiology 38:752-760.

Nieuwenhuis S, Yeung N, van den Wildenberg W, Ridderinkhof KR (2003) Electrophysiological correlates of anterior cingulate function in a go/no-go task: Effects of response conflict and trial type frequency. Cognitive Affective & Behavioral Neuroscience 3:17-26.

Nimchinsky EA, Gilissen E, Allman JM, Perl DP, Erwin JM, Hof PR (1999) A neuronal morphologic type unique to humans and great apes. Proceedings of the National Academy of Sciences of the United States of America 96:5268-5273.

Noppeney U, Price CJ (2002) Retrieval of visual, auditory, and abstract semantics. Neuroimage 15:917-926.

Ochsner KN, Kosslyn SM, Cosgrove GR, Cassem EH, Price BH, Nierenberg AA, Rauch SL (2001) Deficits in visual cognition and attention following bilateral anterior cingulotomy. Neuropsychologia 39:219-230.

O'Doherty J, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ (2004) Dissociable roles of ventral and dorsal striatum in instrumental conditioning. Science 304:452-454.

O'Doherty J, Kringelbach ML, Rolls ET, Hornak J, Andrews C (2001) Abstract reward and punishment representations in the human orbitofrontal cortex. Nature Neuroscience 4:95-102.

O'Doherty JP, Dayan P, Friston K, Critchley H, Dolan RJ (2003) Temporal difference models and reward-related learning in the human brain. Neuron 38:329-337.

O'Doherty JP, Hampton A, Kim H (2007) Model-based fMRI and its application to reward learning and decision making. In: Reward and Decision Making in Corticobasal Ganglia Networks (Balleine BW, Doya K, Odoherty J, Sakagami M, eds), pp 35-53. Oxford: Blackwell Publishing.

Ogawa S, Lee TM, Kay AR, Tank DW (1990) Brain magnetic-resonance-imaging with contrast dependent on blood oxygenation. Proceedings of the National Academy of Sciences of the United States of America 87:9868-9872.

Ongur D, Ferry AT, Price JL (2003) Architectonic subdivision of the human orbital and medial prefrontal cortex. Journal of Comparative Neurology 460:425-449.

Ongur D, Price JL (2000) The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. Cerebral Cortex 10:206-219.

Onishi KH, Bailargeon R (2005) Do 15-month-old infants understand false beliefs? Science 308:255-258.

Padoa-Schioppa C, Assad JA (2006) Neurons in the orbitofrontal cortex encode economic value. Nature 441:223-226.

Padoa-Schioppa C, Assad JA (2008) The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. Nature Neuroscience 11:95-102.

Palmen S, van Engeland H, Hof PR, Schmitz C (2004) Neuropathological findings in autism. Brain 127:2572-2583.

Palomero-Gallagher N, Mohlberg H, Zilles K, Vogt B (2008) Cytology and receptor architecture of human anterior cingulate cortex. Journal of Comparative Neurology 508:906-926.

Pandya DN, Vanhoesen GW, Mesulam MM (1981) Efferent connections of the cingulate gyrus in the rhesus-monkey. Experimental Brain Research 42:319-330.

Passingham RE, Bengtsson SL, Lau HC (2010) Medial frontal cortex: from self-generated action to reflection on one's own performance. Trends in Cognitive Sciences 14:16-21.

Passingham RE, Stephan KE, Kotter R (2002) The anatomical basis of functional localization in the cortex. Nature Reviews Neuroscience 3:606-616.

Paus T (2001) Primate anterior cingulate cortex: Where motor control, drive and cognition interface. Nature Reviews Neuroscience 2:417-424.

Paus T, Tomaiuolo F, Otaky N, MacDonald D, Petrides M, Atlas J, Morris R, Evans AC (1996) Human cingulate and paracingulate sulci: Pattern, variability, asymmetry, and probabilistic map. Cerebral Cortex 6:207-214.

Pearson JM, Heilbronner SR, Barack DL, Hayden BY, Platt ML (2011) Posterior cingulate cortex: adapting behavior to a changing world. Trends in Cognitive Sciences 15:143-151.

Pelphrey KA, Morris JP, McCarthy G (2004) Grasping the intentions of others: The perceived intentionality of an action influences activity in the superior temporal sulcus during social perception. Journal of Cognitive Neuroscience 16:1706-1716.

Perner J, Aichhorn M, Kronbichler M, Staffen W, Ladurner G (2006) Thinking of mental and other representations: The roles of left and right temporo-parietal junction. Social Neuroscience 1:245-258.

Perner J, Lang B (1999) Development of theory of mind and executive control. Trends in Cognitive Sciences 3:337-344.

Peters J, Buchel C (2009) Overlapping and Distinct Neural Systems Code for Subjective Value during Intertemporal and Risky Decision Making. Journal of Neuroscience 29:15727-15734.

Peters J, Buchel C (2010) Episodic Future Thinking Reduces Reward Delay Discounting through an Enhancement of Prefrontal-Mediotemporal Interactions. Neuron 66:138-148.

Petrides M, Pandya DN (1999) Dorsolateral prefrontal cortex: comparative cytoarchitectonic analysis in the human and the macaque brain and corticocortical connection patterns. European Journal of Neuroscience 11:1011-1036.

Petrides M, Pandya DN (2006) Efferent association pathways originating in the caudal prefrontal cortex in the macaque monkey. Journal of Comparative Neurology 498:227-251.

Petrides M, Pandya DN (2007) Efferent association pathways from the rostral prefrontal cortex in the macaque monkey. Journal of Neuroscience 27:11573-11586.

Peyron R, Laurent B, Garcia-Larrea L (2000) Functional imaging of brain responses to pain. A review and meta-analysis (2000). Neurophysiologie Clinique-Clinical Neurophysiology 30:263-288.

Phan KL, Wager T, Taylor SF, Liberzon I (2002) Functional neuroanatomy of emotion: A meta-analysis of emotion activation studies in PET and fMRI. Neuroimage 16:331-348.

Phillips ML, Drevets WC, Rauch SL, Lane R (2003) Neurobiology of emotion perception I: The neural basis of normal emotion perception. Biological Psychiatry 54:504-514.

Phillips PEM, Walton ME, Jhou TC (2007) Calculating utility: preclinical evidence for cost-benefit analysis by mesolimbic dopamine. Psychopharmacology 191:483-495.

Picard N, Strick PL (1996) Motor areas of the medial wall: A review of their location and functional activation. Cerebral Cortex 6:342-353.

Pine A, Seymour B, Roiser JP, Bossaerts P, Friston KJ, Curran HV, Dolan RJ (2009) Encoding of Marginal Utility across Time in the Human Brain. Journal of Neuroscience 29:9575-9581.

Pochon JB, Riis J, Sanfey AG, Nystrom LE, Cohen JD (2008) Functional imaging of decision conflict. Journal of Neuroscience 28:3468-3473.

Pool JL, Ransohoff J (1949) Autonomic effects on stimulating rostral portion of cingulate gyri in man. Journal of neurophysiology 12:385-392.

Posner MI, Rothbart MK (1998) Attention, self-regulation and consciousness. Philosophical Transactions of the Royal Society of London B Biological Sciences 353:1915-1927.

Posner MI, Rothbart MK, Sheese BE, Tang Y (2007) The anterior cingulate gyrus and the mechanism of self-regulation. Cognitive Affective & Behavioral Neuroscience 7:391-395.

Prehn K, Wartenburger I, Meriau K, Scheibe C, Goodenough OR, Villringer A, Van der Meer E, Heekeren HR (2008) Individual differences in moral judgment competence influence neural correlates of socio-normative judgments. Social Cognitive and Affective Neuroscience 3:33-46.

Preuschoff K, Quartz SR, Bossaerts P (2008) Human insula activation reflects risk prediction errors as well as risk. Journal of Neuroscience 28:2745-2752.

Prevost C, Pessiglione M, Metereau E, Clery-Melin ML, Dreher JC (2010) Separate Valuation Subsystems for Delay and Effort Decision Costs. Journal of Neuroscience 30:14080-14090.

Puce A, Allison T, Bentin S, Gore JC, McCarthy G (1998) Temporal cortex activation in humans viewing eye and mouth movements. Journal of Neuroscience 18:2188-2199.

Puce A, Perrett D (2003) Electrophysiology and brain imaging of biological motion. Philosophical Transactions of the Royal Society of London Series B-Biological Sciences 358:435-445.

Quartz SR (2009) Reason, emotion and decision-making: risk and reward computation with feeling. Trends in Cognitive Sciences 13:209-215.

Quilodran R, Rothe M, Procyk E (2008) Behavioral shifts and action valuation in the anterior cingulate cortex. Neuron 57:314-325.

Rachlin H, Raineri A, Cross D (1991) subjective-probability and delay. Journal of the Experimental Analysis of Behavior 55:233-244.

Rainville P, Duncan GH, Price DD, Carrier B, Bushnell MC (1997) Pain affect encoded in human anterior cingulate but not somatosensory cortex. Science 277:968-971.

Ramnani N, Behrens TEJ, Penny W, Matthews PM (2004a) New approaches for exploring anatomical and functional connectivity in the human brain. Biological Psychiatry 56:613-619.

Ramnani N, Elliott R, Athwal BS, Passinghm RE (2004b) Prediction error for free monetary reward in the human prefrontal cortex. Neuroimage 23:777-786.

Ramnani N, Miall RC (2003) Instructed delay activity in the human prefrontal cortex is modulated by monetary reward expectation. Cerebral Cortex 13:318-327.

Ramnani N, Miall RC (2004) A system in the human brain for predicting the actions of others. Nature Neuroscience 7:85-90.

Ramnani N, Owen AM (2004) Anterior prefrontal cortex: Insights into function from anatomy and neuroimaging. Nature Reviews Neuroscience 5:184-194.

Reithler J, van Mier HI, Peters JC, Goebel R (2007) Nonvisual motor learning influences abstract action observation. Current Biology 17:1201-1207.

Rescorla RA, Wagner AR (1972) Classical Conditioning II: Current Research and Theory. In, pp 64–99. New York: Appleton-Century Crofts.

Richards JB, Mitchell SH, DeWit H, Seiden LS (1997) Determination of discount functions in rats with an adjusting-amount procedure. Journal of the Experimental Analysis of Behavior 67:353-366.

Rilling JK, Goldsmith DR, Glenn AL, Jairam MR, Elfenbein HA, Dagenais JE, Murdock CD, Pagnoni G (2008) The neural correlates of the affective response to unreciprocated cooperation. Neuropsychologia 46:1256-1266.

Rilling JK, Sanfey AG, Aronson JA, Nystrom LE, Cohen JD (2004) The neural correlates of theory of mind within interpersonal interactions. Neuroimage 22:1694-1703.

Rizzolatti G, Craighero L (2004) The mirror-neuron system. Annual Review of Neuroscience 27:169-192.

Rizzolatti G, Fadiga L, Gallese V, Fogassi L (1996) Premotor cortex and the recognition of motor actions. Cognitive Brain Research 3:131-141.

Roesch MR, Calu DJ, Burke KA, Schoenbaum G (2007) Should I stay or should I go? Transformation of time-discounted rewards in orbitofrontal cortex and associated brain circuits. In: Reward and Decision Making in Corticobasal Ganglia Networks (Balleine BW, Doya K, Odoherty J, Sakagami M, eds), pp 21-34. Oxford: Blackwell Publishing.

Roesch MR, Taylor AR, Schoenbaum G (2006) Encoding of time-discounted rewards in orbitofrontal cortex is independent of value representation. Neuron 51:509-520.

Rogers RD, Ramnani N, Mackay C, Wilson JL, Jezzard P, Carter CS, Smith SM (2004) Distinct portions of anterior cingulate cortex and medial prefrontal cortex are activated by reward processing in separable phases of decision-making cognition. Biological Psychiatry 55:594-602.

Rolls ET (2000) The orbitofrontal cortex and reward. Cerebral Cortex 10:284-294.

Rolls ET, McCabe C, Redoute J (2008) Expected value, reward outcome, and temporal difference error representations in a probabilistic decision task. Cerebral Cortex 18:652-663.

Rothmayr C, Sodian B, Hajak G, Dohnel K, Meinhardt J, Sommer M (2011) Common and distinct neural networks for false-belief reasoning and inhibitory control. Neuroimage 56:1705-1713.

Rubia K, Russell T, Overmeyer S, Brammer MJ, Bullmore ET, Sharma T, Simmons A, Williams SCR, Giampietro V, Andrew CM, Taylor E (2001) Mapping motor inhibition: Conjunctive brain activations across different versions of go/no-go and stop tasks. Neuroimage 13:250-261.

Rudebeck PH, Bannerman DM, Rushworth MFS (2008) The contribution of distinct subregions of the ventromedial frontal cortex to emotion, social behavior, and decision making. Cognitive Affective & Behavioral Neuroscience 8:485-497.

Rudebeck PH, Behrens TE, Kennerley SW, Baxter MG, Buckley MJ, Walton ME, Rushworth MFS (2008) Frontal Cortex Subregions Play Distinct Roles in Choices between Actions and Stimuli. Journal of Neuroscience 28:13775-13785.

Rudebeck PH, Buckley MJ, Walton ME, Rushworth MFS (2006a) A role for the macaque anterior cingulate gyrus in social valuation. Science 313:1310-1312.

Rudebeck PH, Walton ME, Millette BHP, Shirley E, Rushworth MFS, Bannerman DM (2007) Distinct contributions of frontal areas to emotion and social behaviour in the rat. European Journal of Neuroscience 26:2315-2326.

Rudebeck PH, Walton ME, Smyth AN, Bannerman DM, Rushworth MFS (2006b) Separate neural pathways process different decision costs. Nature Neuroscience 9:1161-1168.

Rushworth MFS, Behrens TEJ (2008) Choice, uncertainty and value in prefrontal and cingulate cortex. Nature Neuroscience 11:389-397.

Rushworth MFS, Behrens TEJ, Rudebeck PH, Walton ME (2007) Contrasting roles for cingulate and orbitofrontal cortex in decisions and social behaviour. Trends in Cognitive Sciences 11:168-176.

Rushworth MFS, Mars RB, Summerfield C (2009) General mechanisms for making decisions? Current Opinion in Neurobiology 19:75-83.

Rushworth MFS, Walton ME, Kennerley SW, Bannerman DM (2004) Action sets and decisions in the medial frontal cortex. Trends in Cognitive Sciences 8:410-417.

Salamone JD, Correa M, Farrar A, Mingote SM (2007) Effort-related functions of nucleus accumbens dopamine and associated forebrain circuits. Psychopharmacology 191:461-482.

Sallet J, Quilodran R, Rothe M, Vezoli J, Joseph JP, Procyk E (2007) Expectations, gains, and losses in the anterior cingulate cortex. Cognitive Affective & Behavioral Neuroscience 7:327-336.

Saltmarsh R, Mitchell P (1999) Indelible evidence of false belief: Confronting young children with video recordings of themselves. Psychologia 42:145-159.

Saltmarsh R, Mitchell P, Robinson E (1995) Realism and childrens early grasp of mental representation - belief-based judgments in the state change task. Cognition 57:297-325.

Sanfey AG, Rilling JK, Aronson JA, Nystrom LE, Cohen JD (2003) The neural basis of economic decision-making in the ultimatum game. Science 300:1755-1758.

Saver JL, Damasio AR (1991) Preserved access and processing of social knowledge in a patient with acquired sociopathy due to ventromedial frontal damage. Neuropsychologia 29:1241-1249.

Saxe R, Kanwisher N (2003) People thinking about thinking people - The role of the temporo-parietal junction in “theory of mind”. Neuroimage 19:1835-1842.

Saxe R, Powell LJ (2006) It's the thought that counts: Specific brain regions for one component of theory of mind. Psychological Science 17:692-699.

Saxe R, Schulz LE, Jiang YV (2006) Reading minds versus following rules: Dissociating theory of mind and executive control in the brain. Social Neuroscience 1:284-298.

Saxe R, Wexler A (2005) Making sense of another mind: The role of the right temporo-parietal junction. Neuropsychologia 43:1391-1399.

Schnell K, Bluschke S, Konradt B, Walter H (2011) Functional relations of empathy and mentalizing: An fMRI study on the neural basis of cognitive empathy. Neuroimage 54:1743-1754.

Schoenbaum G, Roesch MR, Stalnaker TA, Takahashi YK (2009) A new perspective on the role of the orbitofrontal cortex in adaptive behaviour. Nature Reviews Neuroscience 10:885-892.

Schubotz RI (2007) Prediction of external events with our motor system: towards a new framework. Trends in Cognitive Sciences 11:211-218.

Schultz W (1998) Predictive reward signal of dopamine neurons. Journal of Neurophysiology 80:1-27.

Schultz W (2002) Getting formal with dopamine and reward. Neuron 36:241-263.

Schultz W (2006) Behavioral theories and the neurophysiology of reward. Annual Review of Psychology 57:87-115.

Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. Science 275:1593-1599.

Schultz W, Dickinson A (2000) Neuronal coding of prediction errors. Annual Review of Neuroscience 23:473-500.

Schultz W, Tremblay L, Hollerman JR (2000) Reward processing in primate orbitofrontal cortex and basal ganglia. Cerebral Cortex 10:272-283.

Seltzer B, Pandya DN (1989) Frontal-lobe connections of the superior temporal sulcus in the rhesus-monkey. Journal of Comparative Neurology 281:97-113.

Seo H, Lee D (2007) Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. Journal of Neuroscience 27:8366-8377.

Seymour B, Daw N, Dayan P, Singer T, Dolan R (2007) Differential encoding of losses and gains in the human striatum. Journal of Neuroscience 27:4826-4831.

Seymour B, O'Doherty JP, Dayan P, Koltzenburg M, Jones AK, Dolan RJ, Friston KJ, Frackowiak RS (2004) Temporal difference models describe higher-order learning in humans. Nature 429:664-667.

Shackman AJ, Salomons TV, Slagter HA, Fox AS, Winter JJ, Davidson RJ (2011) The integration of negative affect, pain and cognitive control in the cingulate cortex. Nature Reviews Neuroscience 12:154-167.

Shane MS, Stevens M, Harenski CL, Kiehl KA (2008) Neural correlates of the processing of another's mistakes: A possible underpinning for social and observational learning. Neuroimage 42:450-459.

Shane MS, Stevens MC, Harenski CL, Kiehl KA (2009) Double dissociation between perspective-taking and empathic-concern as predictors of hemodynamic response to another's mistakes. Social Cognitive and Affective Neuroscience 4:111-118.

Shidara M, Richmond BJ (2002) Anterior cingulate: Single neuronal signals related to degree of reward expectancy. Science 296:1709-1711.

Showers MJ (1959) The cingulate gyrus: additional motor area and cortical autonomic regulator. The Journal of comparative neurology 112:231-301.

Shultz S, Dunbar RIM (2007) The evolution of the social brain: anthropoid primates contrast with other vertebrates. Proceedings of the Royal Society B-Biological Sciences 274:2429-2436.

Shultz S, Dunbar RIM (2010) Social bonds in birds are associated with brain size and contingent on the correlated evolution of life-history and increased parental investment. Biological Journal of the Linnean Society 100:111-123.

Silk JB (2007) Social components of fitness in primate groups. Science 317:1347-1351.

Silk JB, Alberts SC, Altmann J (2003) Social bonds of female baboons enhance infant survival. Science 302:1231-1234.

Simms ML, Kemper TL, Timbie CM, Bauman ML, Blatt GJ (2009) The anterior cingulate cortex in autism: heterogeneity of qualitative and quantitative cytoarchitectonic features suggests possible subgroups. Acta Neuropathologica 118:673-684.

Singer T, Critchley HD, Preuschoff K (2009) A common role of insula in feelings, empathy and uncertainty. Trends in Cognitive Sciences 13:334-340.

Singer T, Seymour B, O'Doherty J, Kaube H, Dolan RJ, Frith CD (2004) Empathy for pain involves the affective but not sensory components of pain. Science 303:1157-1162.

Smith AT, Singh KD, Balsters JH (2007) A comment on the severity of the effects of non-white noise in fMRI time-series. Neuroimage 36:282-288.

Smith BW, Mitchell DGV, Hardin MG, Jazbec S, Fridberg D, Blair RJR, Ernst M (2009) Neural substrates of reward magnitude, probability, and risk during a wheel of fortune decision-making task. Neuroimage 44:600-609.

Smith DV, Hayden BY, Truong TK, Song AW, Platt ML, Huettel SA (2010) Distinct Value Signals in Anterior and Posterior Ventromedial Prefrontal Cortex. Journal of Neuroscience 30:2490-2495.

Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TEJ, Johansen-Berg H, Bannister PR, De Luca M, Drobnjak I, Flitney DE, Niazy RK, Saunders J, Vickers J, Zhang YY, De Stefano N, Brady JM, Matthews PM (2004) Advances in functional and structural MR image analysis and implementation as FSL. Neuroimage 23:S208-S219.

Sohn MH, Albert MV, Jung KJ, Carter CS, Anderson JR (2007) Anticipation of conflict monitoring in the anterior cingulate cortex and the prefrontal cortex. Proceedings of the National Academy of Sciences of the United States of America 104:10330-10334.

Sommer M, Dohnel K, Sodian B, Meinhardt J, Thoermer C, Hajak G (2007) Neural correlates of true and false belief reasoning. Neuroimage 35:1378-1384.

Sommer M, Meinhardt J, Eichenmuller K, Sodian B, Dohnel K, Hajak G (2010) Modulation of the cortical false belief network during development. Brain Research 1354:123-131.

Southgate V, Senju A, Csibra G (2007) Action anticipation through attribution of false belief by 2-year-olds. Psychological Science 18:587-592.

Spitzer M, Fischbacher U, Herrnberger B, Gron G, Fehr E (2007) The neural signature of social norm compliance. Neuron 56:185-196.

Spooren W, LyndBalta E, Mitchell S, Haber SN (1996) Ventral pallidostriatal pathway in the monkey: Evidence for modulation of basal ganglia circuits. Journal of Comparative Neurology 370:295-312.

Stehling MK, Turner R, Mansfield P (1991) Echo-planar imaging - magnetic-resonance-imaging in a fraction of a second. Science 254:43-50.

Stenger VA, Boada FE, Noll DC (2000) Three-dimensional tailored RF pulses for the reduction of susceptibility artifacts in T-2*-weighted functional MRI. Magnetic Resonance in Medicine 44:525-531.

Stephens DW, Charnov EL (1982) Optimal foraging - some simple stochastic-models. Behavioral Ecology and Sociobiology 10:251-263.

Stroop JR (1935) Studies of interference in serial verbal reactions. Journal of Experimental Psychology:643-662.

Summerfield C, Egner T, Greene M, Koechlin E, Mangels J, Hirsch J (2006) Predictive codes for forthcoming perception in the frontal cortex. Science 314:1311-1314.

Summerfield C, Koechlin E (2008) A neural representation of prior information during perceptual inference. Neuron 59:336-347.

Surian L, Caldi S, Sperber D (2007) Attribution of beliefs by 13-month-old infants. Psychological Science 18:580-586.

Sutton RS, Barto AG (1981) Toward a modern theory of adaptive networks - expectation and prediction. Psychological Review 88:135-170.

Sutton RS, Barto AG (1998) Reinforcement learning: an introduction. Cambridge, Massachusetts: MIT press.

Takada M, Tokuno H, Hamada I, Inase M, Ito Y, Imanishi M, Hasegawa N, Akazawa T, Hatanaka N, Nambu A (2001) Organization of inputs from cingulate motor areas to basal ganglia in macaque monkey. European Journal of Neuroscience 14:1633-1650.

Terreberry RR, Neafsey EJ (1983) Rat medial frontal-cortex - a visceral motor region with a direct projection to the solitary nucleus. Brain Research 278:245-249.

Terreberry RR, Neafsey EJ (1984) The effects of medial prefrontal cortex stimulation on heart rate in the awake rat. Society for Neuroscience Abstracts 10:614.

Terreberry RR, Neafsey EJ (1987) The rat medial frontal-cortex projects directly to autonomic regions of the brain-stem. Brain Research Bulletin 19:639-649.

Thaler D, Chen YC, Nixon PD, Stern CE, Passingham RE (1995) the functions of the medial premotor cortex .1. simple learned movements. Experimental Brain Research 102:445-460.

Thompson JC, Hardee JE, Panayiotou A, Crewther D, Puce A (2007) Common and distinct brain activation to viewing dynamic sequences of face and hand movements. Neuroimage 37:966-973.

Thornton A, McAuliffe K (2006) Teaching in wild meerkats. Science 313:227-229.

Tobler PN, Fiorillo CD, Schultz W (2005) Adaptive coding of reward value by dopamine neurons. Science 307:1642-1645.

Tomlin D, Kayali MA, King-Casas B, Anen C, Camerer CF, Quartz SR, Montague PR (2006) Agent-specific responses in the cingulate cortex during economic exchanges. Science 312:1047-1050.

Tranel D, Bechara A, Denburg NL (2002) Asymmetric functional roles of right and left ventromedial prefrontal cortices in social conduct, decision-making, and emotional processing. Cortex 38:589-612.

Tricomi E, Rangel A, Camerer CF, O'Doherty JP (2010) Neural evidence for inequality-averse social preferences. Nature 463:1089-U1109.

Ullen F, Forssberg H, Ehrsson HH (2003) Neural networks for the coordination of the hands in time. Journal of Neurophysiology 89:1126-1135.

Van Overwalle F, Baetens K (2009) Understanding others' actions and goals by mirror and mentalizing systems: A meta-analysis. Neuroimage 48:564-584.

van Schie HT, Mars RB, Coles MGH, Bekkering H (2004) Modulation of activity in medial frontal and motor cortices during error observation. Nature Neuroscience 7:549-554.

Vianna EPM, Van Hoesen G, Parvizi J (2002) Efferent connections of the primate posterior cingulate, retrosplenial, and mesial parietal cortices. Society for Neuroscience Abstract.

Vogeley K, Bussfeld P, Newen A, Herrmann S, Happe F, Falkai P, Maier W, Shah NJ, Fink GR, Zilles K (2001) Mind reading: Neural mechanisms of theory of mind and self-perspective. Neuroimage 14:170-181.

Vogt BA, Nimchinsky EA, Vogt LJ, Hof PR (1995) Human cingulate cortex - surface-features, flat maps, and cytoarchitecture. Journal of Comparative Neurology 359:490-506.

Vogt BA, Pandya DN (1987) Cingulate cortex of the rhesus-monkey .2. cortical afferents. Journal of Comparative Neurology 262:271-289.

Vogt BA, Pandya DN, Rosene DL (1987) Cingulate cortex of the rhesus-monkey .1. cytoarchitecture and thalamic afferents. Journal of Comparative Neurology 262:256-270.

Waelti P, Dickinson A, Schultz W (2001) Dopamine responses comply with basic assumptions of formal learning theory. Nature 412:43-48.

Wager TD, Rilling JK, Smith EE, Sokolik A, Casey KL, Davidson RJ, Kosslyn SM, Rose RM, Cohen JD (2004) Placebo-induced changes in fMRI in the anticipation and experience of pain. Science 303:1162-1167.

Walter H, Adenzato M, Ciaramidaro A, Enrici I, Pia L, Bara BG (2004) Understanding intentions in social interaction: The role of the anterior paracingulate cortex. Journal of Cognitive Neuroscience 16:1854-1863.

Walter H, Ciaramidaro A, Adenzato M, Vasic N, Ardito RB, Erk S, Bara BG (2009) Dysfunction of the social brain in schizophrenia is modulated by intention type: An fMRI study. Social Cognitive and Affective Neuroscience 4:166-176.

Walton ME, Devlin JT, Rushworth MFS (2004) Interactions between decision making and performance monitoring within prefrontal cortex. Nature Neuroscience 7:1259-1265.

Walton ME, Kennerley SW, Bannerman DM, Phillips PEM, Rushworth MFS (2006) Weighing up the benefits of work: Behavioral and neural analyses of effort-related decision making. Neural Networks 19:1302-1314.

Walton ME, Mars RB (2007) Probing human and monkey anterior cingulate cortex in variable environments. Cognitive Affective & Behavioral Neuroscience 7:413-422.

Wanat MJ, Kuhnen CM, Phillips PEM (2010) Delays Conferred by Escalating Costs Modulate Dopamine Release to Rewards But Not Their Predictors. Journal of Neuroscience 30:12020-12027.

Wang Y, Shima K, Sawamura H, Tanji J (2001) Spatial distribution of cingulate cells projecting to the primary, supplementary, and pre-supplementary motor areas: a retrograde multiple labeling study in the macaque monkey. Neuroscience Research 39:39-49.

Ward AA, Jr. (1948) The cingular gyrus, area 24. Journal of neurophysiology 11:13-23.

Watson KK, Jones TK, Allman JM (2006) Dendritic architecture of the von Economo neurons. Neuroscience 141:1107-1112.

Wellman HM, Cross D, Watson J (2001) Meta-analysis of theory-of-mind development: The truth about false belief. Child Development 72:655-684.

Whalen PJ, Bush G, Shin LM, Rauch SL (2006) The emotional counting Stroop: a task for assessing emotional interference during brain imaging. Nature Protocols 1:293-296.

Wilkinson HA, Davidson KM, Davidson RI (1999) Bilateral anterior cingulotomy for chronic noncancer pain. Neurosurgery 45:1129-1134.

Williams JHG, Whiten A, Suddendorf T, Perrett DI (2001) Imitation, mirror neurons and autism. Neuroscience and Biobehavioral Reviews 25:287-295.

Williams SM, Goldman-Rakic PS (1998) Widespread origin of the primate mesofrontal dopamine system. Cerebral Cortex 8:321-345.

Williams ZM, Bush G, Rauch SL, Cosgrove GR, Eskandar EN (2004) Human anterior cingulate neurons and the integration of monetary reward with motor responses. Nature Neuroscience 7:1370-1375.

Wimmer H, Perner J (1983) Beliefs about beliefs - representation and constraining function of wrong beliefs in young childrens understanding of deception. Cognition 13:103-128.

Woolrich MW, Ripley BD, Brady M, Smith SM (2001) Temporal autocorrelation in univariate linear modeling of FMRI data. Neuroimage 14:1370-1386.

Wunderlich K, Rangel A, O'Doherty JP (2009) Neural computations underlying action-based decision making in the human brain. Proceedings of the National Academy of Sciences of the United States of America 106:17199-17204.

Yeterian EH, Pandya DN (1991) Prefrontostriatal connections in relation to cortical architectonic organization in rhesus-monkeys. Journal of Comparative Neurology 312:43-67.

Yeterian EH, Pandya DN (1993) Striatal connections of the parietal association cortices in rhesus-monkeys. Journal of Comparative Neurology 332:175-197

Yeung N, Holroyd CB, Cohen JD (2005) ERP correlates of feedback and reward processing in the presence and absence of response choice. Cerebral Cortex 15:535-544.

Yeung N, Nieuwenhuis S (2009) Dissociating Response Conflict and Error Likelihood in Anterior Cingulate Cortex. Journal of Neuroscience 29:14506-14510.

Yoshida W, Dolan RJ, Friston KJ (2008) Game Theory of Mind. Plos Computational Biology 4.

Young L, Camprodon JA, Hauser M, Pascual-Leone A, Saxe R (2010b) Disruption of the right temporoparietal junction with transcranial magnetic stimulation reduces the role of beliefs in moral judgments. Proceedings of the National Academy of Sciences of the United States of America 107:6753-6758.

Young L, Cushman F, Hauser M, Saxe R (2007) The neural basis of the interaction between theory of mind and moral judgment. Proceedings of the National Academy of Sciences of the United States of America 104:8235-8240.

Young L, Dodell-Feder D, Saxe R (2010a) What gets the attention of the temporo-parietal junction? An fMRI investigation of attention and theory of mind. Neuropsychologia 48:2658-2664.

Young L, Saxe R (2009) An fMRI Investigation of Spontaneous Mental State Inference for Moral Judgment. Journal of Cognitive Neuroscience 21:1396-1405.

Zaitchik D, Walker C, Miller S, LaViolette P, Feczko E, Dickerson BC (2010) Mental state attribution and the temporoparietal junction: An fMRI study comparing belief, emotion, and perception. Neuropsychologia 48:2528-2536.

Appendix

Appendix A (1) – Experiment one debrief sheet

Debriefing Sheet.

Thank you for taking part in the study. We would like to provide you with some further information about the experiment so that you are fully aware of the methods and aims.

The aim of the project was to investigate neural responses that occur when participants are monitoring the predictions others' make about an outcome. We are particularly interested in when the predictions made by another person are false. Earlier we told you that on some trials you were observing the decisions made by the 'decision-maker'. However, these responses were not related to another person, but were computer generated. Also, the decision-maker was not another participant, but a confederate (one of the researchers).

It was necessary to use these measures because they allowed us to maintain control over the experiment whilst maintaining your belief that you were observing the responses and predictions related to another person's outcomes. In light of this new information, we would like to give you another opportunity to consider whether or not you will permit us to use your data. You are free to withdraw consent for us to use your data without giving reasons.

We would be grateful if you could indicate your decision on the consent form below.

Second Consent form

ID number…….

Monitoring others' false predictions: An fMRI study.

You have been asked to participate in a study about the brain processes that occur when monitoring others' predictions, which is being carried out by Dr. Narender Ramnani, Matthew Apps, Vivien Ainley and Robin Green.

Have you (please circle yes or no):

- Read the information sheet about the study? yes no
- Had an opportunity to ask questions? yes no
- Got satisfactory answers to your questions? yes no
- Understood that you're free to withdraw from the study at any time, without giving a reason yes no

(and without it affecting your care/ education if applicable)?

Do you give consent to use your data in light of the new information provided after the study?

yes no

Signature_____ Name in block letters _____Date _____

NB: This consent form will be stored separately from the anonymous information you provide.

Appendix A (2) – Experiment two debrief sheet

Thank you for taking part in the study. We would like to provide you with some further information about the experiment so that you are fully aware of the methods and aims.

The aim if the project was to investigate neural responses that occur when participants learn that someone else has made a correct or incorrect response and when they are providing them with feedback. Earlier we told you that you were observing the decisions made by the 'Learner'. However, these responses were not related to another person, but were computer generated. Also, the Learner was not another participant, but a confederate (one of the researchers).

It was necessary to use these measures because they allowed us to maintain control over the experiment whilst maintaining your belief that you were observing the decisions made by another person. In light of this new information, we would like to give you another opportunity to consider whether or not you will permit us to use your data. You are free to withdraw consent for us to use your data without giving reasons.

We would be grateful if you could indicate your decision on the consent form below.

……

Second Consent form

ID number…….

Teaching Other's: An fMRI study.

You have been asked to participate in a study about the brain processes that occur when monitoring others' predictions, which is being carried out by Dr. Narender Ramnani and Matthew Apps.

Have you (please circle yes or no):

- Read the information sheet about the study? yes no
- Had an opportunity to ask questions? yes no
- Got satisfactory answers to your questions? yes no
- Understood that you're free to withdraw from the study at any time, without giving a reason yes no

(and without it affecting your care/ education if applicable)?

Do you give consent to use your data in light of the new information provided after the study?

yes no

Signature_____ Name in block letters_____Date _____

NB: This consent form will be stored separately from the anonymous information you provide.

Appendix A (3) – Experiment three debrief sheet

Thank you for taking part in the study. We would like to provide you with some further information about the experiment so that you are fully aware of the methods and aims.

The aim of the project was to investigate neural responses that occur when you have to expend different amounts of effort for differing financial rewards. We are particularly interested in examining what happens when you are thinking about the amount of effort another is expending. We previously told you that you were performing this task with a training partner situated outside the scanner. However, in fact, your training partner is a member of the research team. The responses you saw on the screen on each trial were actually pre-programmed computer-controlled responses. It was necessary to use these measures because they allowed us to maintain control over the experiment and examine how individuals understand the effort expended by another more reliably.

In light of this new information, we would like to give you another opportunity to consider whether or not you will permit us to use your data. You are free to withdraw consent for us to use your data without giving reasons.

We would be grateful if you could indicate your decision on the consent form below.

Second Consent form

ID number…….

Effort and Rewards: An fMRI study.

You have participated in a study about the brain processes that occur when you put in different amounts of effort to receive financial rewards,which is being carried out by Matthew Apps, Erman Misirlisoy, Arleta Woznica and Dr. Narender Ramnani.

Have you (please circle yes or no):

- Read the information sheet about the study? yes no
- Had an opportunity to ask questions? yes no
- Got satisfactory answers to your questions? yes no
- Understood that you're free to withdraw from the study at any time, without giving a reason yes no

(and without it affecting your care/ education if applicable)?

Do you give consent to use your data in light of the new information provided after the study?
yes no

Signature_____ Name in block letters_____Date _____

NB: This consent form will be stored separately from the anonymous information you provide.

Appendix A (4) – Experiment four debrief sheet

Thank you for taking part in the study. We would like to provide you with some further information about the experiment so that you are fully aware of the methods and aims.

The aim of the project was to investigate neural responses that occur when people make financial decisions. We are particularly interested in examining what happens when participants learn about how the majority of people make these financial decisions. We previously told you that we had collected data from 102 participants and established that 69% of people agreed on the better option. We also told you that on each trial you received this information. However, in fact, we have not collected data from any participants other than those involved in this MRI experiment. The information you received was not that of any other individual, but was part of the experimental design. The experiment was designed so that you may not always choose the same option as the majority of other people. It was necessary to use these measures because they allowed us to maintain control over the experiment and investigate how people learn about how others make financial decisions.

We would like to stress that the way people perform this task is highly variable and there is no abnormal behaviour. The task does not therefore indicate that the way you make financial decisions is any different from how other people make them. In light of this new information, we would like to give you another opportunity to consider whether or not you will permit us to use your data. You are free to withdraw consent for us to use your data without giving reasons.

We would be grateful if you could indicate your decision on the consent form below.

Second Consent form

ID number…….

Financial Decision-Making: An fMRI study.

You have been asked to participate in a study about the brain processes that occur when making financial decisions, which is being carried out by Dr. Narender Ramnani and Matthew Apps.

Have you (please circle yes or no):

- Read the information sheet about the study? yes no
- Had an opportunity to ask questions? yes no
- Got satisfactory answers to your questions? yes no
- Understood that you're free to withdraw from the study at any time, without giving a reason yes no

  (and without it affecting your care/ education if applicable)?

Do you give consent to use your data in light of the new information provided after the study? yes no

Signature_____ Name in block letters
_____Date _____

NB: This consent form will be stored separately from the anonymous information you provide.

Appendix B (1) – Quality assurance for chapter three.

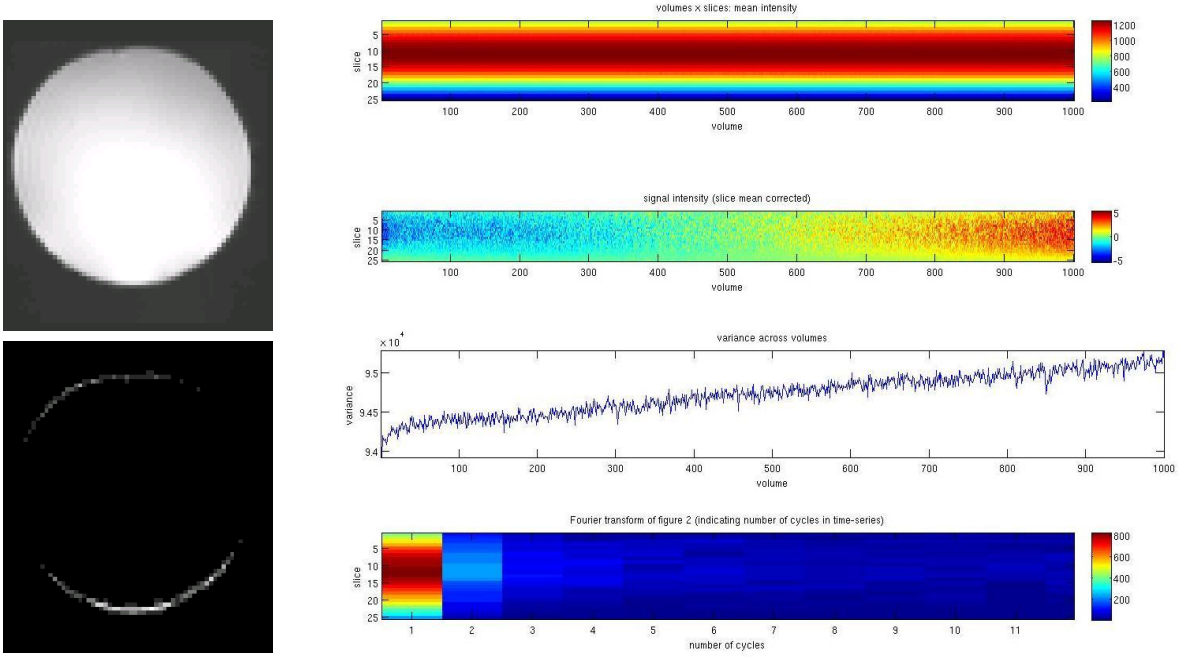

The results displayed are from "phantom" data collected one week prior to onset of collection of data for chapter three. In the top left is the mean image of the timeseries. Bottom left is a standard deviation image across the timeseries. On the right hand side, from the top, are the mean intensity for each slice, the slice mean corrected signal intensity, the variance across volumes and a fast fourier transform on each slice. Data collection for this experiment was collected over three months, two more Q.A.'s were also performed during this time. The parameters used were the same as those used during the experimental scan.

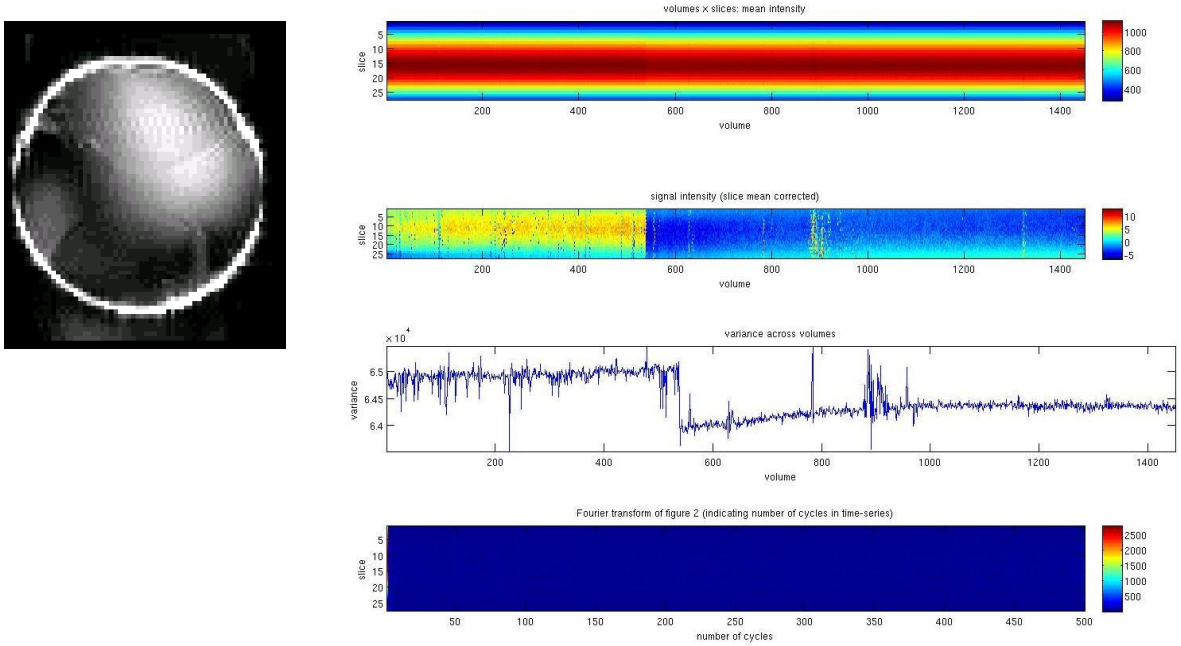

Illustration of a Q.A analysis where significant problems were identified prior to chapter three data collection. This data was collected three months prior to that displayed above. On the left is the standard deviation image, which shows considerable ghosting and time-varying changes in signal intensity in the centre of the phantom. On the right are the same plots as in the figure above.The images on the right show the considerable variability in signal that occurred across the whole volume.

Appendix B (2) – Quality assurance for chapter four, five and six

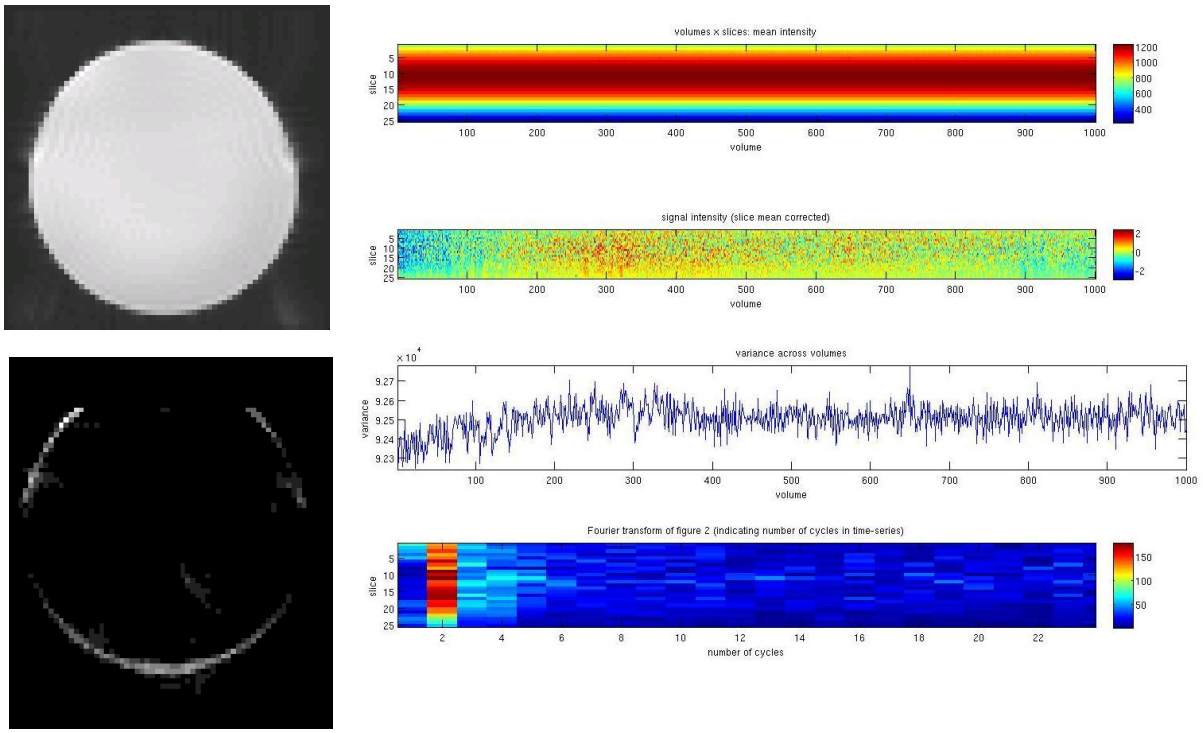

Experiment two Q.A. results displayed are from "phantom" data collected two weeks before data collection for chapter four began. The displayed images are the same as outlined for experiment one above.

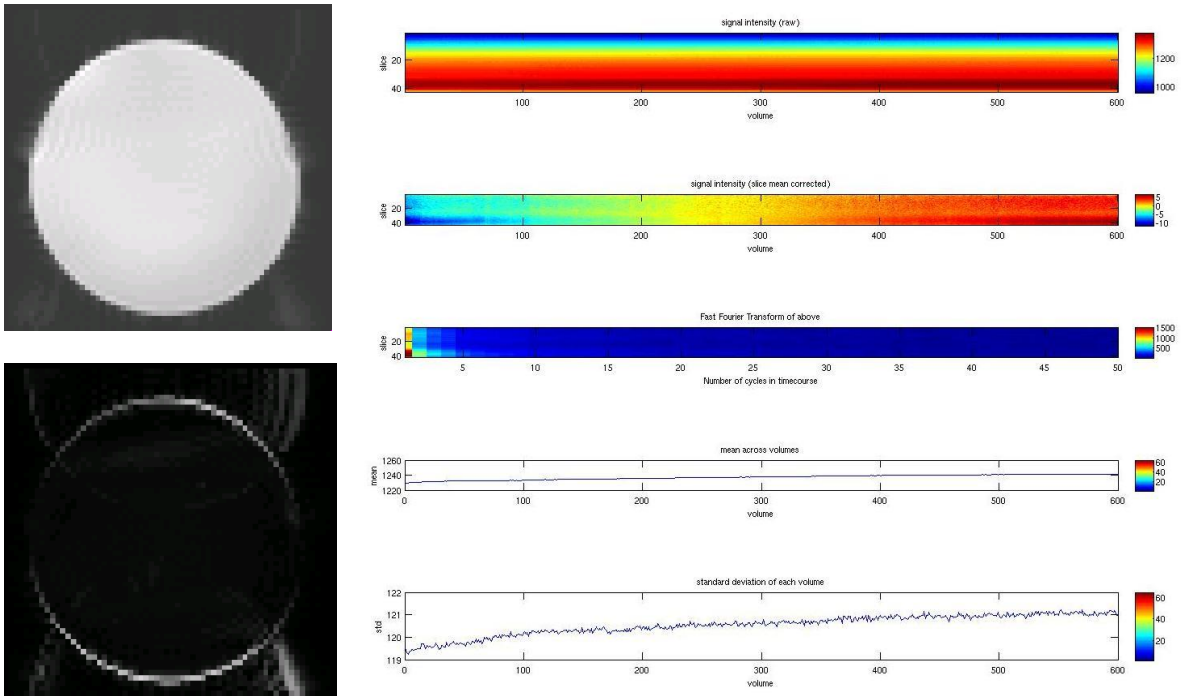

Experiment three and four Q.A. results displated are from "phantom" data collected one week before chapter six data collection began and five weeks before chapter five data collection began. Data collection was conducted sequentially for these experiments. In the top left is the mean image of the timeseries. Bottom left is a standard deviation image across the timeseries. On the right hand side, from the top, are the mean intensity for each slice, the slice mean corrected signal intensity, a fourier transform on each slice, the mean across volumes and the standard deviation across volumes. Data collection was carried out over 10 weeks for both experiments. An additional QA was conducted 6 weeks

Appendix C (1) – Questions about belief in deception for chapter three

- How did you find the task?
- How did [confederate's name] perform on the task?
- Did you feel it was easy to distinguish between the trials of the Computer and that those of your training partner?
- Where you thinking about the predicted outcome on the trials differently from the Computer?
- Was the performance of your training partner surprising or was you would have expected?

---------- DEBRIEF SHEET GIVEN TO SUBJECT ----------

- Are you surprised at what you have just read?
- Did you believe that the behaviour you observed was that of the training partner during the experiment?
- We would be grateful if you did not inform anyone about the nature of the deception in the experiment.

Questions about belief in deception for chapter four

- How did you find the task?
- How did [confederate's name] perform on the task?
- When they were making a response, were you thinking specifically about them making correct or incorrect responses or were you just thinking more generally about it being a correct or incorrect response?
- Was the performance of your training partner surprising or was you would have expected?

---------- DEBRIEF SHEET GIVEN TO SUBJECT ----------

- Are you surprised at what you have just read?
- Did you believe that the behaviour you observed was that of the training partner during the experiment?
- We would be grateful if you did not inform anyone about the nature of the deception in the experiment.

Appendix C (2) – Questions about belief in deception for chapter five

- How did you find the task?
- How did [confederate's name] perform on the task?
- Could you provide a rough estimate of how much each of you has earned?
- Did you find the performance of the effort task easy?
- Did you find it easy to make judgements about how much money they would earn?

---------- DEBRIEF SHEET GIVEN TO SUBJECT ----------

- Are you surprised at what you have just read?
- Did you believe that the behaviour you observed was that of the training partner during the experiment?
- Did you believe that they were earning money in the same way that you were?
- We would be grateful if you did not inform anyone about the nature of the deception in the experiment.

Questions about belief in deception for chapter six

- How did you find the task?
- Did you find it easy to learn the group majority behaviour?
- Were you similar to this majority or different?
- Were you surprised at the group behaviour?

---------- DEBRIEF SHEET GIVEN TO SUBJECT ----------

- Are you surprised at what you have just read?
- Did you believe that the behaviour you learnt really represented what the majority of other people did on the task?
- Did you believe that we performed an experiment on over 100 subjects before this fMRI experiment?
- We would be grateful if you did not inform anyone about the nature of the deception in the experiment.

Appendix D – Peristimulus time histogram plots from chapter five.

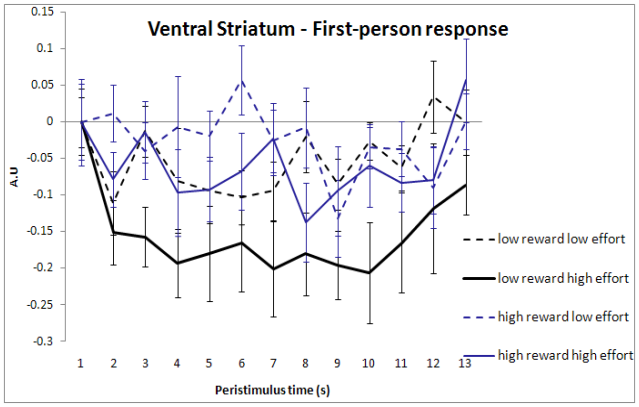

PSTH plot for activity in the peak ventral striatum voxel shown in figure 5.5

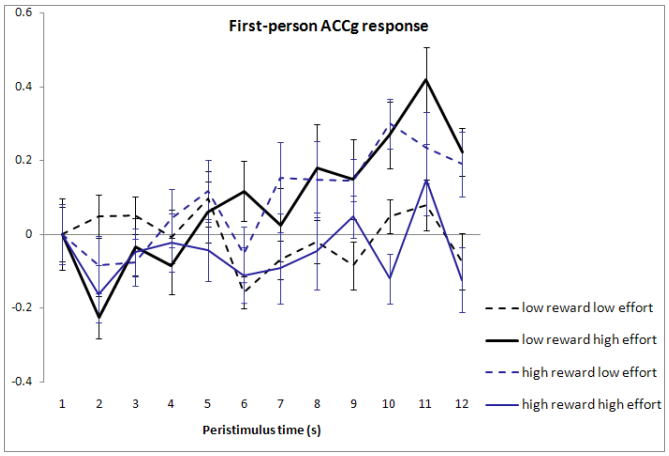

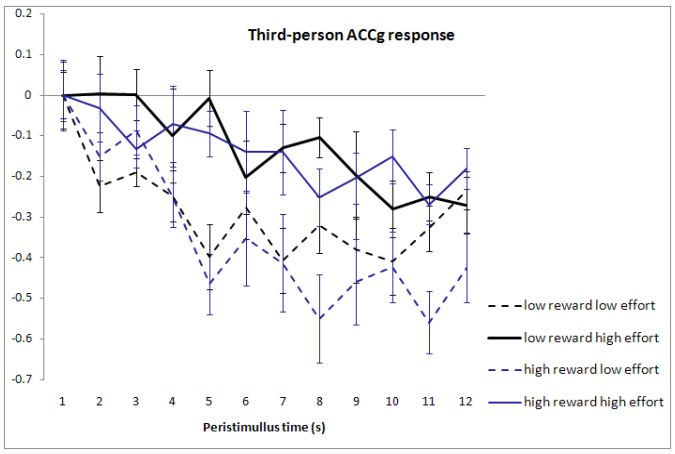

PSTH plots for activity in the peak ACCg voxel shown in figure 5.6. the top shows the response to the first-person instruction cues and the bottom shows the response to the third-person instruction cues

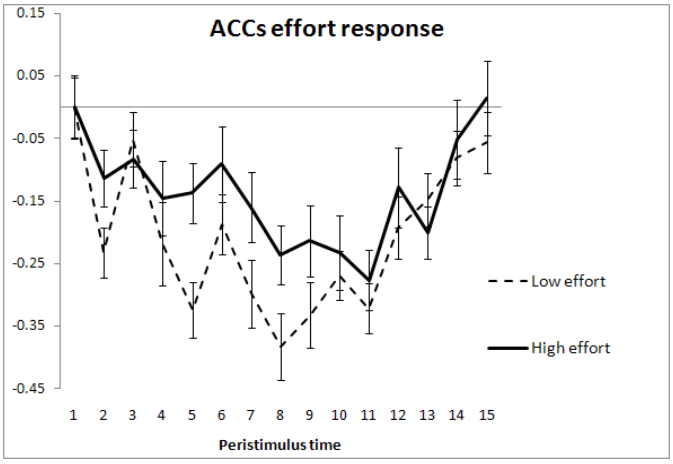

PSTH plot for activity in the peak ACCs voxel shown in figure 5.8. For display purposes the responses were collapsed across the two levels of Reward and Agency, to give the overall response for High and Low Effort.

Appendix E – Coordinates used for small volume correction in chapters five and six

Anatomical region	*MNI Coordinate*
Ventral Striatum	-6, 10 -6 12, 8, -2
Putamen	-26, -8, -2 20, 14, -12
Cerebellum	10, -46, -28
Midbrain	-6, -20, -8 4, -20, -8

Coordinates used for small volume correction in chapter 5, taken from Croxson et al., (2009). These regions responded to first-person net value in that experiment. Spheres with an 8mm radius around these coordinates were made into a mask that was used for the small volume correction; this mask was then added to the ACCs mask for correction of any of contrasts for the first-person conditions. Activity was only found in the ventral striatum in chapter 5.

Anatomical region	*MNI Coordinate*
Ventral Striatum	-11, 5, 1
Posterior Cingulate/retropslenial cortex	-3, -43, 37

Coordinates used for small volume correction in chapter 6, taken from Kable and Glimcher (2007). Activity in these regions varied with the subjective value of temporally discounted rewards. This mask was then added to the ACCs mask for correction of the subjective value parameter in chapter 6.

Appendix F(1) – masks used for small volume corrections in chapters three and four

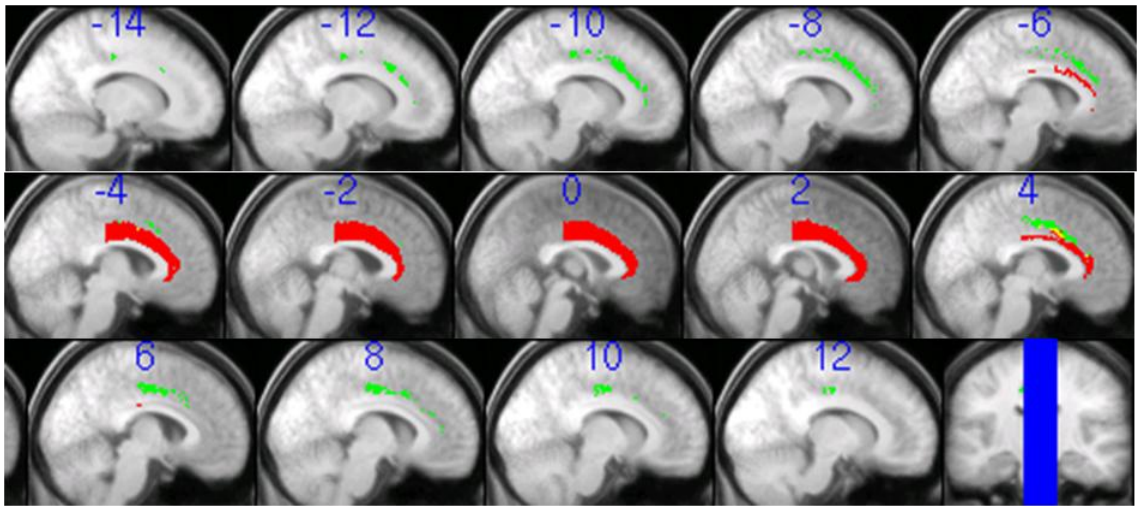

Chapter 3 masks. The ACCs (green) and ACCg (red) masks used for small volume correction in chapter three. Voxels which are yellow overlap between the two masks. The masks contained only voxels which were present in 80% of subjects scanned in this study and overlaid on the mean anatomical image.

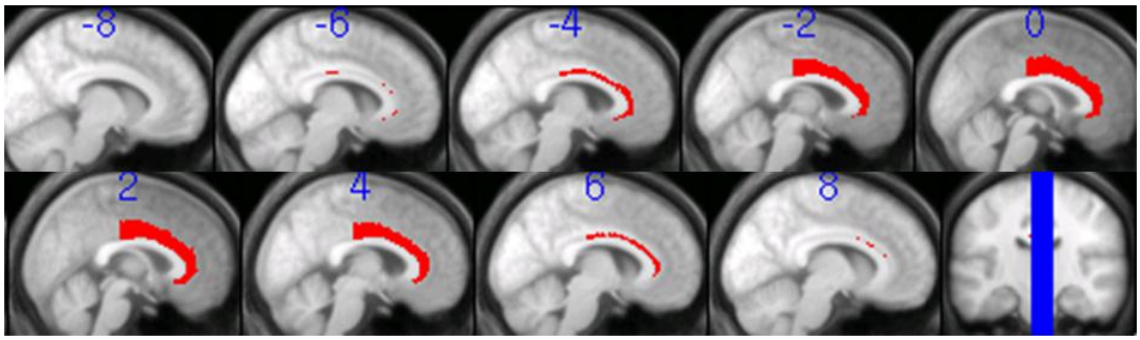

Chapter 4 mask.The ACCg (red) masks used for small volume correction in chapter four. No hypothesis was tested for the ACCs and so no mask was created. The mask contains only voxels which were present in 80% of subjects and overlaid on the mean anatomical image.

Appendix F(2) – masks used for small volume corrections in chapter five

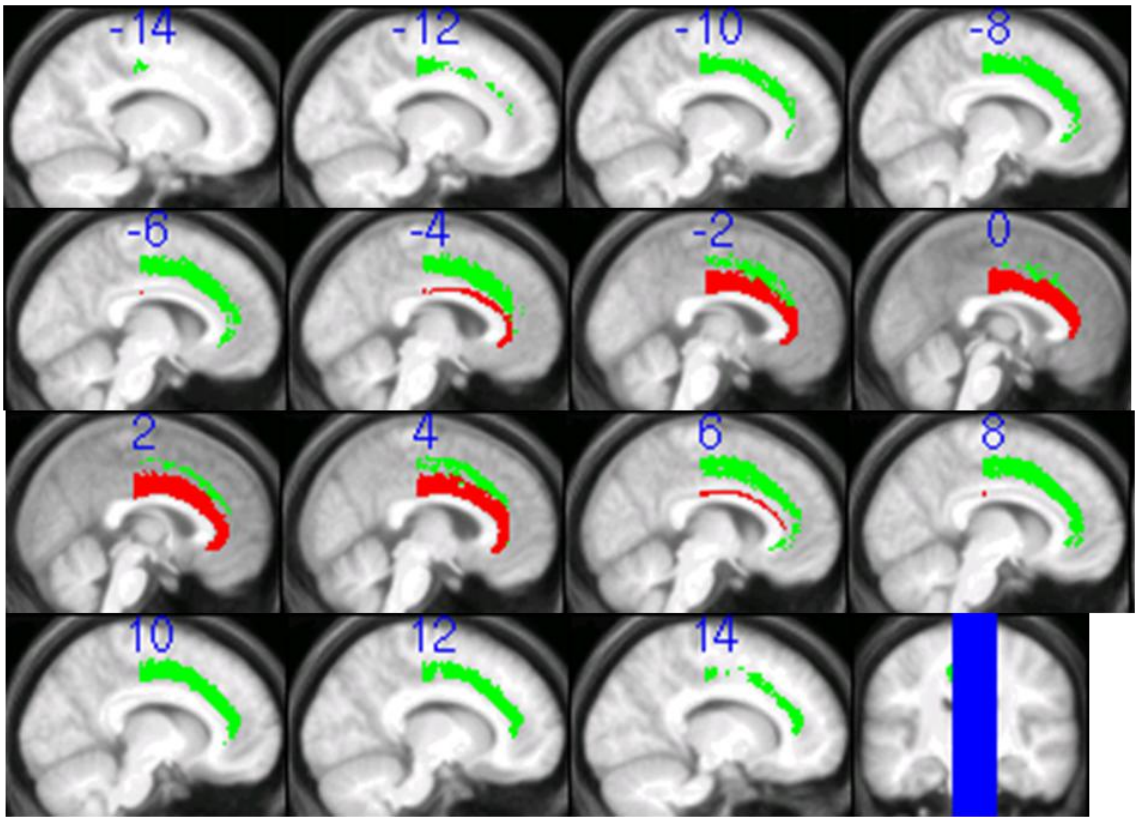

Chapter 5 masks.The ACCs (green) and ACCg (red) masks used for small volume correction in chapter five. Voxels which are yellow overlap between the two masks. The masks contained only voxels which were present in 80% of subjects scanned in this study and overlaid on the mean anatomical image.

Appendix F(3) – masks used for small volume corrections in chapter six

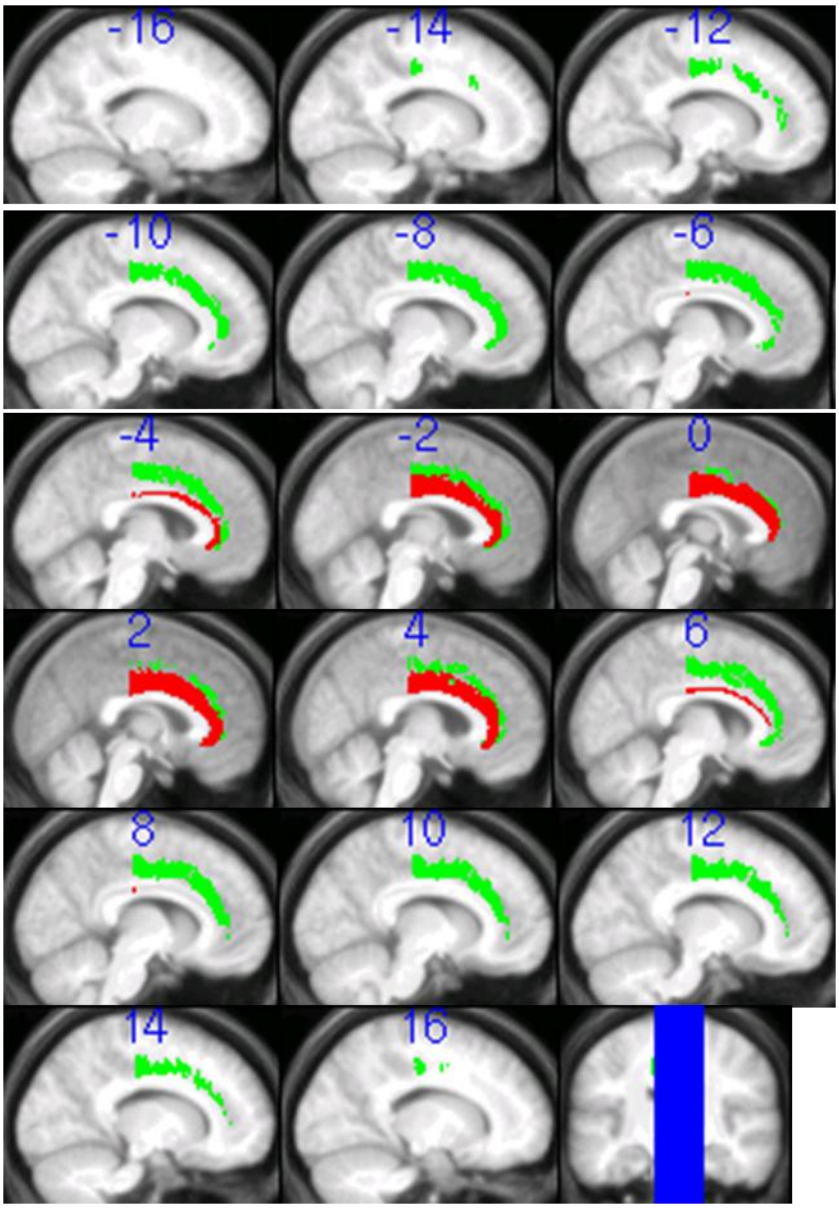

Chapter 6 masks.The ACCs (green) and ACCg (red) masks used for small volume correction in chapter six. Voxels which are yellow overlap between the two masks. The masks contained only voxels which were present in 80% of subjects scanned in this study and overlaid on the mean anatomical image.